

Robust Coreset Construction for Distributed Machine Learning

Hanlin Lu, *Student Member, IEEE*, Ming-Ju Li, Ting He, *Senior Member, IEEE*, Shiqiang Wang, *Member, IEEE*, Vijaykrishnan Narayanan, *Fellow, IEEE*, and Kevin S Chan, *Senior Member, IEEE*

Abstract—Motivated by the need of solving machine learning problems over distributed datasets, we explore the use of *coreset* to reduce the communication overhead. Coreset is a summary of the original dataset in the form of a small weighted set in the same sample space. Compared to other data summaries, coreset has the advantage that it can be used as a proxy of the original dataset, potentially for different applications. However, existing coreset construction algorithms are each tailor-made for a specific machine learning problem. Thus, to solve different machine learning problems, one has to collect coresets of different types, defeating the purpose of saving communication overhead. We resolve this dilemma by developing coreset construction algorithms based on k-means/median clustering, that give a guaranteed approximation for a broad range of machine learning problems with sufficiently continuous cost functions. Through evaluations on diverse datasets and machine learning problems, we verify the robust performance of the proposed algorithms.

Index Terms—Coreset, distributed machine learning, distributed k-means, distributed k-median.

I. INTRODUCTION

The recent decade has observed a dramatic growth in distributed data generation, powered by various Internet of Things (IoT) applications and social networking applications. It has been predicted that the rate of such distributed data generation will exceed the current Internet capacity in the near future [2]. This phenomenon presents both opportunities and challenges for machine learning. On the one hand, the real-time and location-based nature of the distributed data enables novel applications based on machine learning, such as augmented reality and cognitive assistance. On the other hand, the distributed nature of the data sources, coupled with the difficulty of collecting the data to a central location due to network bandwidth, energy, and/or privacy constraints, calls for a fundamentally different way of applying machine learning that is suitable for distributed datasets.

Broadly speaking, three approaches have been proposed for distributed machine learning: *sharing the output*, *sharing the model*, and *sharing the data*. In the first approach [3], data sources independently compute and share outputs of local




Figure 1. Application scenario (ML i : machine learning model i).

models, which are then aggregated into a global output (e.g., by majority vote). In the second approach [4], [5], [6], data sources share models learned on local data, which are then aggregated into a global model (e.g., by taking weighted average). In the third approach [7], [8], [9], data sources share summaries of their local data, which are then used to compute a global model. Each approach has its pros and cons: the first approach usually has the smallest communication overhead, but only supports one query; the second approach builds a global model that can be used for multiple queries, but only supports one model; the third approach may incur a larger communication overhead, but can potentially support multiple models, hence amortizing the overhead.

In this work, we take the third approach, with a particular interest in supporting diverse machine learning models. Figure 1 illustrates a typical application scenario in the context of *mobile edge computing* [5], where data sources report local summaries to an edge server, which then computes various models from these summaries.

In particular, we consider data summarization using *coreset* [10]. A coreset is a small weighted dataset as a proxy of the original dataset with *provable approximation guarantees*. Compared to other data summaries (e.g., sketches), a coreset preserves the sample space of the original dataset, and is hence more convenient to use, e.g., a classifier learned from the coreset can classify a new data point, the principle components learned from the coreset can be used for feature selection, both *in the original sample space*. Algorithms have been developed to construct coresets for various machine learning problems (see Section II-B), such that the model learned on the coreset approximates the model learned on the original dataset. However, existing coreset construction algorithms are *tailor-made* for specific machine learning problems, which means that we have to collect different coresets to solve different problems. The question we raise is: *Is there a coreset that is good for a broad set of machine learning problems?*

In this work, we answer the above question affirmatively

H. Lu, M. Li, T. He, and V. Narayanan are with Pennsylvania State University, University Park, PA 16802 USA (email: {hxl263, mxl592, tzh58, vxn9}@psu.edu).

S. Wang is with IBM T. J. Watson Center, Yorktown Heights, NY 10598 USA (email: wangshiq@us.ibm.com).

K. Chan is with Army Research Laboratory, Adelphi, MD 20783 USA (email: kevin.s.chan.civ@mail.mil).

A preliminary version of this work was presented at GLOBECOM'19 [1].

This research was partly sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. Narayanan and Hanlin were partly supported by NSF 1317560. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

by proving that a particular type of coreset, generated by k -means/median clustering, can give a good approximation for a broad set of machine learning problems.

A. Related Work

Distributed learning is considered one of the most promising lines of research for large-scale learning [3], particularly for naturally distributed data. The main challenge in distributed learning is to incorporate information from each distributed dataset, without the high overhead of collecting all the data.

Traditionally, this is achieved by collecting the outputs of learned models or the models themselves [11]. The first approach (i.e., collecting outputs) is more popular among earlier works. For example, [12] proposed various heuristic decision rules (e.g., majority vote) to combine outputs of local classifiers, and [13] proposed to train a global classifier using labeled outputs of local classifiers. The solution in [13] was modified in [14] to improve efficiency for large-scale distributed data, and extended in [11] to include various ways of composing the global training set. The idea was later used to build a descriptive model from distributed data [15]. To further improve the accuracy, a distributed-pasting-votes framework was proposed in [16] to learn sets of classifiers (ensembles).

The second approach (i.e., collecting models) is more useful when we want to learn not just one answer, but the rule to give answers. For example, the distributed boosting framework in [17] requires nodes to share locally trained classifiers, and the federated learning framework in [4], [18] requires nodes to report locally learned models to a single node, which then aggregates the models and broadcasts the result to others.

Meanwhile, research on data summarization has inspired a third approach: collecting data summaries. Data summaries, e.g., coresets, sketches, projections [19], [20], [21], are derived datasets that are much smaller than the original dataset, and can hence be transferred to a central location with a low communication overhead. This approach has been adopted in recent works, e.g., [7], [8], [9], [22]. We are particularly interested in a specific type of data summary, *coreset*, as it can be used as a proxy of the original dataset. See Section II-B for a detailed review of related works on coreset.

B. Summary of Contributions

We are the first to explore using coreset to support diverse machine learning problems on distributed data. Specifically:

- 1) We empirically show that although existing coreset construction algorithms are designed for specific machine learning problems, an algorithm based on k -means clustering yields good performance for different problems.
- 2) We harden the above observation by proving that the optimal k -clustering (including k -means/median) gives a coreset that provides a guaranteed approximation for any machine learning problem with a sufficiently continuous cost function (Theorem III.1). We further prove that the same holds for the coreset given by a suboptimal k -clustering algorithm, as long as it satisfies certain assumptions (Theorem III.2).

- 3) We adapt an existing algorithm designed to support distributed k -clustering to construct a robust coreset over distributed data with a very low communication overhead.
- 4) Our evaluations on diverse machine learning problems and datasets verify that k -clustering (especially k -means) and its approximations provide coresets good for learning a variety of machine learning models.

Roadmap. Section II reviews the background on coreset. Section III presents our main theoretical results on the universal performance guarantee of k -clustering-based coreset. Section IV presents our distributed coreset construction algorithm. Section V evaluates the proposed algorithm. Section VI concludes the paper.

II. BACKGROUND

A. Coreset and Machine Learning

Many machine learning problems can be cast as a cost (or loss) minimization problem. Given a dataset in d -dimensional space $P \subseteq \mathbb{R}^d$, a generic machine learning problem over P can be characterized by a solution space \mathcal{X} , a *per-point cost function* $\text{cost}(p, x)$ ($p \in P, x \in \mathcal{X}$), and an *overall cost function* $\text{cost}(P, x)$ ($x \in \mathcal{X}$) that aggregates the per-point costs over P . For generality, we consider P to be a weighted set, where each $p \in P$ has weight w_p . Let $w_{\min} := \min_{p \in P} w_p$ denote the minimum weight. For an unweighted dataset, we have $w_p \equiv 1$. The machine learning problem is then to solve

$$x^* = \arg \min_{x \in \mathcal{X}} \text{cost}(P, x) \quad (1)$$

for the optimal model parameter x^* .

Example: Let $\text{dist}(p, x) := \|p - x\|_2$ denote the Euclidean distance between points p and x . The *minimum enclosing ball (MEB)* problem [10] aims at minimizing the maximum distance between any data point and a center, i.e., $\text{cost}(p, x) = \text{dist}(p, x)$, $\text{cost}(P, x) = \max_{p \in P} \text{cost}(p, x)$, and $\mathcal{X} = \mathbb{R}^d$. The *k -means clustering* problem aims at minimizing the weighted sum of the squared distance between each data point and the nearest center in a set of k centers, i.e., $\text{cost}(p, x) = \min_{x_i \in \mathcal{X}} \text{dist}(p, x_i)^2$, $\text{cost}(P, x) = \sum_{p \in P} w_p \text{cost}(p, x)$, and $\mathcal{X} = \{x := \{x_i\}_{i=1}^k : x_i \in \mathbb{R}^d\}$.

Typically, the overall cost is defined as: (i) *sum cost*, i.e., $\text{cost}(P, x) = \sum_{p \in P} w_p \text{cost}(p, x)$ (e.g., k -means), or (ii) *maximum cost*, i.e., $\text{cost}(P, x) = \max_{p \in P} \text{cost}(p, x)$ (e.g., MEB).

A coreset is a small weighted dataset in the same space as the original dataset that approximates the original dataset in terms of cost, formally defined below.

Definition II.1 ([23]). A weighted set $S \subseteq \mathbb{R}^d$ with weights u_q ($q \in S$) is an ϵ -coreset for P with respect to (w.r.t.) $\text{cost}(P, x)$ ($x \in \mathcal{X}$) if $\forall x \in \mathcal{X}$,

$$(1 - \epsilon) \text{cost}(P, x) \leq \text{cost}(S, x) \leq (1 + \epsilon) \text{cost}(P, x), \quad (2)$$

where $\text{cost}(S, x)$ is defined in the same way as $\text{cost}(P, x)$, i.e., $\text{cost}(S, x) = \sum_{q \in S} u_q \text{cost}(q, x)$ for *sum cost*, and $\text{cost}(S, x) = \max_{q \in S} \text{cost}(q, x)$ for *maximum cost*.

From Definition II.1, it is clear that the quality of a cores­et depends on the cost function it needs to approximate, and hence the machine learning problem it supports.

B. Cores­et Construction Algorithms

Because of the dependence on the cost function (Definition II.1), existing cores­et construction algorithms are tailor-made for specific machine learning problems. Here we briefly summarize common approaches for cores­et construction and representative algorithms, and refer to [19], [20] for detailed surveys.

1) *Farthest point algorithms*: Originally proposed for MEB [10], [24], these algorithms iteratively add to the cores­et a point far enough or farthest from the current center, and stop when the enclosing ball of the cores­et, expanded by $1 + \epsilon$, includes all data points. This cores­et has been used to compute ϵ -approximation to several clustering problems, including k -center clustering, 1-cylinder clustering, and k -flat clustering [10], [25]. As *support vector machine (SVM)* training can be formulated as MEB problems [26], similar algorithms have been used to support SVM [26], [27]. Variations have been used for dimensionality reduction [28] and probabilistic MEB [29]. These algorithms are considered as variations of the Frank-Wolfe algorithm [30].

2) *Random sampling algorithms*: These algorithms construct a cores­et by sampling from the original dataset. The basic version, uniform sampling, usually requires a large cores­et size to achieve a good approximation. Advanced versions use *sensitivity sampling* [31], where each data point is sampled with a probability proportional to its contribution to the overall cost. Proposed for numerical integration [31], the idea was extended into a framework supporting projective clustering problems that include k -median/means and *principle component analysis (PCA)* as special cases [23]. The framework has been used to generate cores­ets for other problems, e.g., dictionary learning [32] and dependency networks [33], and is further generalized in [34]. Although the framework can instantiate algorithms for different machine learning problems by plugging in different cost functions, the resulting cores­et only guarantees approximation for the specific problem defined by the plugged-in cost function.

3) *Geometric decomposition algorithms*: These algorithms divide the sample space or input dataset into partitions, and then selecting points to represent each partition. Specific instances have been developed for weighted facility problems [35], Euclidean graph problems [36], k -means/median [9], [37].

While there are a few works not fully covered by the above approaches, e.g., SVD-based algorithms in [38], [39], the above represents the key approaches used by existing cores­et construction algorithms. Using a generic merge-and-reduce approach in [40], all these algorithms can be used to construct cores­ets of distributed datasets. Of course, the resulting cores­ets are still tailor-made for specific problems. In contrast, we seek cores­et construction algorithms which can construct cores­ets that simultaneously support multiple machine learning problems.

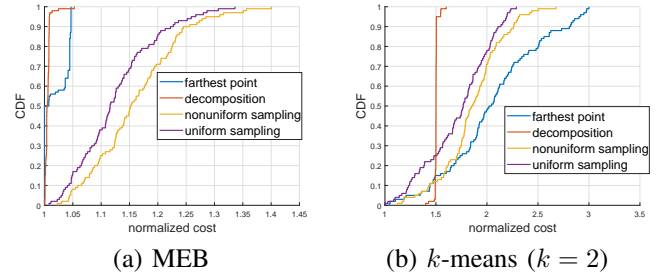


Figure 2. Comparison of cores­et construction algorithms (cores­et size: 8).

III. ROBUST CORESET CONSTRUCTION

Our main result is that selecting *representative points* using clustering techniques yields a cores­et that achieves a good approximation for a broad set of machine learning problems. We will start with a centralized setting in this section, where the raw data reside at a single data source (that needs to compute and report the cores­et to a server as illustrated in Figure 1), and leave the distributed setting where the raw data are distributed across multiple data sources to Section IV.

A. Motivating Experiment

We start with an initial experiment that compares selected algorithms representing the three approaches in Section II-B: (i) the algorithm in [24] (‘farthest point’) representing farthest point algorithms, (ii) the framework in [23] instantiated for k -means (‘nonuniform sampling’) and the uniform sampling algorithm (‘uniform sampling’), both representing random sampling algorithms, and (iii) the algorithm in [9] (‘decomposition’) representing geometric decomposition algorithms. As the algorithm in [24] was designed for MEB and the algorithms in [23], [9] were designed for k -means, we evaluate both MEB and k -means.

The evaluation is based on a synthetic dataset containing 4000 points uniformly distributed in $[1, 50]^3$; evaluations on real datasets will be presented later (Section V). All the algorithms are tuned to have the same average cores­et size. We evaluate the performance of a cores­et S by the *normalized cost*, defined as $\text{cost}(P, x_S) / \text{cost}(P, x^*)$, where x^* is the model learned from the original dataset P , and x_S is the model learned from the cores­et. The smaller the normalized cost, the better the performance. As these cores­et construction algorithms are randomized, we plot the CDF of the normalized costs computed over 100 Monte Carlo runs in Figure 2.

Not surprisingly, a cores­et construction algorithm designed for one problem can perform poorly for another, e.g., the farthest point algorithm [24] designed for MEB performs poorly for k -means. Interestingly, the decomposition algorithm [9], although designed for k -means, also performs well for MEB. This observation triggers a question: Is the superior performance of the decomposition algorithm [9] just a coincidence, or is there something fundamental?

B. The k -clustering Problem

At the core, the decomposition algorithm in [9] constructs a k -point cores­et by partitioning the dataset into k clusters

using k -means clustering, and then using the cluster centers as the coreset points. To analyze its performance in supporting a general machine learning problem, we introduce a few definitions.

Given a weighted dataset $P \subseteq \mathbb{R}^d$ with weight w_p ($p \in P$), and a set $Q = \{q_1, \dots, q_k\}$ of $k \geq 1$ points in \mathbb{R}^d (referred to as *centers*), the cost of clustering P into Q is defined as

$$c(P, Q) = \sum_{p \in P} w_p (\min_{q \in Q} \text{dist}(p, q))^z, \quad (3)$$

for a constant $z > 0$. The k -clustering problem is to find the set of k centers that minimizes (3). For $z = 1$, this is the k -median problem. For $z = 2$, this is the k -means problem. We will use the solution to the k -clustering problem to construct coresets, based on which we can solve general machine learning problems. We use $c(P, \cdot)$ to denote the cost function of this auxiliary problem and $\text{cost}(P, \cdot)$ to denote the cost function of a general machine learning problem of interest.

We denote by $\mu(P)$ the optimal center for 1-clustering of P . It is known that for $z = 2$, $\mu(P)$ is the sample mean:

$$\mu(P) = \frac{1}{\sum_{p \in P} w_p} \sum_{p \in P} w_p \cdot p. \quad (4)$$

We denote by $\text{opt}(P, k)$ the optimal cost for k -clustering of P . It is known that k -means and k -median are both NP-hard problems [41], [42], for which efficient heuristics exist (e.g., Lloyd's algorithm and variations) [43]. Let $\text{approx}(P, k)$ denote the cost of a generic k -clustering algorithm, which always satisfies $\text{approx}(P, k) \geq \text{opt}(P, k)$.

Each set of k centers $Q = \{q_i\}_{i=1}^k$ induces a partition of P into $\{P_1, \dots, P_k\}$, where P_i is the subset of points in P whose closest center in Q is q_i (ties broken arbitrarily). For ease of presentation, we use¹ $\{P_i\}_{i \in [k]}$ to denote the partition induced by the optimal k -clustering, and $\{\tilde{P}_i\}_{i \in [k]}$ to denote the partition induced by a suboptimal k -clustering.

C. Coreset by Optimal k -clustering

We will show that the superior performance of the algorithm in [9] observed in Section III-A is not a coincidence; instead, it is a fundamental property of any coreset computed by k -clustering, as long as the cost function of the machine learning problem satisfies certain continuity conditions.

Sketch of analysis: At a high level, our analysis is based on the following observations:

- 1) If doubling the number of centers only reduces the optimal k -clustering cost by a little, then using two centers instead of one in any cluster gives little reduction to its clustering cost (Lemma III.1).
- 2) If selecting two centers in a cluster P_i gives little reduction to its clustering cost, then all the points in P_i must be close to its center $\mu(P_i)$ (Lemma III.2), as otherwise selecting an outlier as the second center would have reduced the cost substantially.

- 3) If each data point is represented by a coreset point with a similar per-point cost, then the coreset gives a good approximation of the overall cost (Lemmas III.3 and III.4).

Therefore, for any machine learning problem with a sufficiently continuous cost function, if the condition in item (1) is satisfied, then the per-point cost of each k -clustering center will closely approximate the per-point costs of all the points in its cluster, and hence the set of k -clustering centers will give a good coreset (Theorem III.1).

Complete analysis: We now present the precise statements, supported by proofs in Appendix A.

Lemma III.1. *For any $\epsilon' > 0$, if $\text{opt}(P, k) - \text{opt}(P, 2k) \leq \epsilon'$, then $\text{opt}(P_i, 1) - \text{opt}(P_i, 2) \leq \epsilon'$ ($\forall i \in [k]$), where $\{P_i\}_{i=1}^k$ is the partition of P generated by the optimal k -clustering.*

Lemma III.2. *If $\text{opt}(P_i, 1) - \text{opt}(P_i, 2) \leq \epsilon'$, then $\text{dist}(p, \mu(P_i)) \leq (\frac{\epsilon'}{w_{\min}})^{\frac{1}{z}}$, $\forall p \in P_i$.*

Lemma III.3. *For any machine learning problem with cost function $\text{cost}(P, x) = \sum_{p \in P} w_p \text{cost}(p, x)$, if \exists a partition $\{P_i\}_{i=1}^k$ of P such that $\forall x \in \mathcal{X}$, $i \in [k]$, and $p \in P_i$,*

$$(1 - \epsilon) \text{cost}(p, x) \leq \text{cost}(\mu(P_i), x) \leq (1 + \epsilon) \text{cost}(p, x), \quad (5)$$

then $S = \{\mu(P_i)\}_{i=1}^k$ with weight $u_{\mu(P_i)} = \sum_{p \in P_i} w_p$ is an ϵ -coreset for P w.r.t. $\text{cost}(P, x)$.

Lemma III.4. *For any machine learning problem with cost function $\text{cost}(P, x) = \max_{p \in P} \text{cost}(p, x)$, if \exists a partition $\{P_i\}_{i=1}^k$ of P such that (5) holds for any $x \in \mathcal{X}$, $i \in [k]$, and $p \in P_i$, then $S = \{\mu(P_i)\}_{i=1}^k$ (with arbitrary weights) is an ϵ -coreset for P w.r.t. $\text{cost}(P, x)$.*

We now prove the main theorem based on Lemmas III.1–III.4.

Theorem III.1. *If $\text{opt}(P, k) - \text{opt}(P, 2k) \leq w_{\min}(\frac{\epsilon}{\rho})^z$, then the optimal k -clustering of P gives an ϵ -coreset for P w.r.t. both the sum cost and the maximum cost for any per-point cost function satisfying (i) $\text{cost}(p, x) \geq 1$, and (ii) $\text{cost}(p, x)$ is ρ -Lipschitz-continuous in p , $\forall x \in \mathcal{X}$.*

Proof. By Lemma III.1, $\text{opt}(P, k) - \text{opt}(P, 2k) \leq \epsilon'$ implies $\text{opt}(P_i, 1) - \text{opt}(P_i, 2) \leq \epsilon'$, \forall cluster P_i generated by the optimal k -clustering. By Lemma III.2, this in turn implies that $\text{dist}(p, \mu(P_i)) \leq (\frac{\epsilon'}{w_{\min}})^{\frac{1}{z}}$, $\forall p \in P_i$. Because $\text{cost}(p, x)$ is ρ -Lipschitz-continuous in p for all $x \in \mathcal{X}$, we have

$$|\text{cost}(p, x) - \text{cost}(\mu(P_i), x)| \leq \rho(\frac{\epsilon'}{w_{\min}})^{\frac{1}{z}}, \forall x \in \mathcal{X}, p \in P_i. \quad (6)$$

Moreover, as $\text{cost}(p, x) \geq 1$,

$$\frac{|\text{cost}(p, x) - \text{cost}(\mu(P_i), x)|}{\text{cost}(p, x)} \leq \rho(\frac{\epsilon'}{w_{\min}})^{\frac{1}{z}} = \epsilon$$

for $\epsilon' = w_{\min}(\frac{\epsilon}{\rho})^z$. By Lemma III.3, k -clustering gives an ϵ -coreset for P w.r.t. the sum cost; by Lemma III.4, k -clustering gives an ϵ -coreset for P w.r.t. the maximum cost. \square

Often in practice, the coreset size must satisfy some upper bound specified by the maximum communication overhead. In

¹Throughout the paper, for $k \in \mathbb{Z}^+$, $[k] := \{1, \dots, k\}$.

this case, we can rephrase Theorem III.1 to characterize the quality of approximation as a function of the coreset size.

Corollary III.1.1. *Given a maximum coreset size $k \in \mathbb{Z}^+$ (positive integers), for any cost function satisfying the conditions in Theorem III.1, the optimal k -clustering gives an ϵ -coreset for P w.r.t. this cost function, where*

$$\epsilon = \rho \left(\frac{\text{opt}(P, k) - \text{opt}(P, 2k)}{w_{\min}} \right)^{\frac{1}{z}}. \quad (7)$$

Proof. This is a direct implication of Theorem III.1, as setting ϵ by (7) satisfies the condition in Theorem III.1. \square

Remark: Condition (i) in Theorem III.1 is easily satisfied by any machine learning problem with nonnegative per-point costs, as we can add ‘+1’ to the cost function without changing the optimal solution. Even without this condition, a similar proof will show that the coreset S given by k -clustering approximates the original dataset P in that $|\text{cost}(P, x) - \text{cost}(S, x)| \leq \tilde{\epsilon}$ ($\forall x \in \mathcal{X}$), where $\tilde{\epsilon} = \epsilon \sum_{p \in P} w_p$ for the sum cost, and $\tilde{\epsilon} = \epsilon$ for the maximum cost.

Condition (ii) is satisfied by many machine learning problems with distance-based cost functions. For example, for MEB, $\text{cost}(p, x) = \text{dist}(p, x)$, where $x \in \mathbb{R}^d$ denotes the center of the enclosing ball. For any data points $p, p' \in \mathbb{R}^d$, by the triangle inequality, we have:

$$|\text{dist}(p, x) - \text{dist}(p', x)| \leq \text{dist}(p, p'). \quad (8)$$

Hence, its cost function is 1-Lipschitz-continuous (i.e., $\rho = 1$). See Appendix B for more examples. In Section V, we will stress-test our coreset when this condition is violated.

From the proof of Theorem III.1, it is easy to see that the theorem holds as long as the distance between each data point and its nearest k -clustering center is bounded by ϵ/ρ , i.e., $\text{dist}(p, \mu(P_i)) \leq \epsilon/\rho$ for all $i \in [k]$ and $p \in P_i$. This implies that Corollary III.1.1 actually holds for

$$\epsilon = \rho \left(\max_{i \in [k]} \max_{p \in P_i} \text{dist}(p, \mu(P_i)) \right), \quad (9)$$

which can be much smaller than (7) for large datasets.

D. Coreset by Suboptimal k -clustering

While Theorem III.1 and Corollary III.1.1 suggest that the optimal k -clustering gives a good coreset, the k -clustering problem is NP-hard [41], [42]. The question is whether similar performance guarantee holds for the coreset computed by an efficient but suboptimal k -clustering algorithm. To this end, we introduce a few assumptions on the k -clustering algorithm:

Assumption 1 (local optimality): Given the partition $\{\tilde{P}_i\}_{i=1}^k$ generated by the algorithm, the center it selects in each \tilde{P}_i is $\mu(\tilde{P}_i)$, i.e., the optimal 1-clustering center for \tilde{P}_i .

Assumption 2 (self-consistency): For any P and any $k \geq 1$, the cost of the algorithm satisfies

$$\text{approx}(P, 2k) \leq \sum_{i=1}^k \text{approx}(\tilde{P}_i, 2). \quad (10)$$

Assumption 3 (greedy dominance): For any P , the 2-clustering cost of the algorithm satisfies

$$\text{approx}(P, 2) \leq c(P, \{\mu(P), p^*\}), \quad (11)$$

where $c(P, Q)$ is defined in (3), and $p^* := \arg \max_{p \in P} w_p \cdot \text{dist}(p, \mu(P))^z$ is the point with the highest 1-clustering cost.

These are mild assumptions that should be satisfied or approximately satisfied by any good k -clustering algorithm. Assumption 1 is easy to satisfy, as computing the 1-mean is easy (i.e., sample mean), and there exists an algorithm [44] that can compute the 1-median to arbitrary precision in nearly linear time. Assumption 2 means that applying the algorithm for $2k$ -clustering of P should perform no worse than first using the algorithm to partition P into k clusters, and then computing 2-clustering of each cluster. Assumption 3 means that for $k = 2$, the algorithm should perform no worse than a greedy heuristic that starts with the 1-clustering center, and then adds the point with the highest clustering cost as the second center. We will discuss how to ensure these assumptions for the proposed algorithm in Section III-E.

We show that for any k -clustering algorithm satisfying these assumptions, statements analogous to Lemma III.1 and Lemma III.2 can be made (proofs in Appendix A). Let $\{\tilde{P}_i\}_{i=1}^k$ denote the partition of P generated by the k -clustering algorithm.

Lemma III.5. *For any $\epsilon' > 0$, if $\text{approx}(P, k) - \text{approx}(P, 2k) \leq \epsilon'$, then $\text{approx}(\tilde{P}_i, 1) - \text{approx}(\tilde{P}_i, 2) \leq \epsilon'$ for any $i \in [k]$.*

Lemma III.6. *If $\text{approx}(\tilde{P}_i, 1) - \text{approx}(\tilde{P}_i, 2) \leq \epsilon'$, then $\text{dist}(p, \mu(\tilde{P}_i)) \leq (\frac{\epsilon'}{w_{\min}})^{\frac{1}{z}}, \forall p \in \tilde{P}_i$.*

Theorem III.2. *If $\text{approx}(P, k) - \text{approx}(P, 2k) \leq w_{\min}(\frac{\epsilon}{\rho})^z$, where $\text{approx}(P, k)$ is the cost of a (possibly suboptimal) k -clustering algorithm satisfying Assumptions 1–3, then the centers computed by the algorithm for k -clustering of P give an ϵ -coreset for P w.r.t. both the sum cost and the maximum cost for any per-point cost function satisfying (i–ii) in Theorem III.1.*

Proof. The proof follows the same steps as that of Theorem III.1, except that Lemma III.1 is replaced by Lemma III.5, and Lemma III.2 is replaced by Lemma III.6. Note that Lemmas III.3 and III.4 hold for any partition of P , which in this case is $\{\tilde{P}_i\}_{i=1}^k$ generated by the k -clustering algorithm. \square

Similar to Corollary III.1.1, we can rephrase Theorem III.2 to characterize the quality of a coreset of a specified size.

Corollary III.2.1. *Given a maximum coreset size $k \in \mathbb{Z}^+$, for any cost function satisfying the conditions in Theorem III.1 and any k -clustering algorithm satisfying Assumptions 1–3, the centers computed by the algorithm for k -clustering of P give an ϵ -coreset for P w.r.t. the given cost function, where*

$$\epsilon = \rho \left(\frac{\text{approx}(P, k) - \text{approx}(P, 2k)}{w_{\min}} \right)^{\frac{1}{z}}. \quad (12)$$

E. Coreset Construction Algorithm

Based on Theorem III.2, we propose a centralized k -clustering-based coreset construction algorithm, called *Robust*

Algorithm 1: Robust Coreset Construction (P, ϵ, ρ)

input : A weighted set P with minimum weight w_{\min} , approximation error $\epsilon > 0$, Lipschitz constant ρ

output: An ϵ -coreset S for P w.r.t. a cost function satisfying Theorem III.2

```

1 foreach  $k = 1, \dots, |P|$  do
2   if  $\text{approx}(P, k) - \text{approx}(P, 2k) \leq w_{\min}(\frac{\epsilon}{\rho})^z$  then
3     break;
4    $(\{\mu(\tilde{P}_i)\}_{i=1}^k, \{\tilde{P}_i\}_{i=1}^k) \leftarrow k\text{-clustering}(P, k);$ 
5    $S \leftarrow \{\mu(\tilde{P}_i)\}_{i=1}^k$ , where  $\mu(\tilde{P}_i)$  has weight  $\sum_{p \in \tilde{P}_i} w_p$ ;
6 return  $S$ ;
```

Coreset Construction (RCC) (Algorithm 1), which uses a k -clustering algorithm as subroutine in lines 2 and 4. If the coreset size k is predetermined, we can directly start from line 4. The constant $z = 1$ if the adopted clustering algorithm is for k -median, or $z = 2$ if it is for k -means.

The k -clustering subroutine: Algorithm 1 can use any k -clustering algorithm as subroutine, although our performance guarantee holds only if the algorithm satisfies Assumptions 1–3. We note that these assumptions are easy to satisfy if $z = 2$. Consider the standard k -means algorithm (i.e., Lloyd’s algorithm), which iteratively assigns each point to the nearest center and updates the centers to the means of the clusters. Clearly, this algorithm satisfies Assumption 1. Moreover, with the following initialization, it also satisfies Assumptions 2 and 3. For $(2k)$ -clustering of P :

- 1) if $k = 1$, then use the mean $\mu(P)$ and the point p^* with the highest clustering cost as defined in (11) as the initial centers, which helps to satisfy Assumption 3;
- 2) if $k > 1$, then first compute k -clustering of P , and then compute 2-clustering of each of the k clusters (both by calling the same algorithm recursively); finally, use the union of the 2-clustering centers as the initial centers, which helps to satisfy Assumption 2.

Any odd number of initial centers are chosen randomly. Since iterations can only reduce the cost, Lloyd’s algorithm with this initialization satisfies Assumptions 1–3.

In theory, the above initialization plus a Lloyd-style algorithm can satisfy Assumptions 1–3 for an arbitrary $z > 0$, given a subroutine to compute the optimal 1-clustering center $\mu(P)$. For $z = 1$, there is an algorithm to compute $\mu(P)$ to an arbitrary precision in nearly linear time [44].

IV. DISTRIBUTED CORESET CONSTRUCTION

In a distributed setting, the entire dataset P is distributed across n ($n > 1$) nodes (i.e., data sources) v_1, \dots, v_n , where each v_j has a subset $P_j \subseteq P$. We have shown in Section III that the k -clustering centers of P form a robust coreset. However, computing the global k -clustering centers of a distributed dataset is highly non-trivial. Note that a naive solution that only includes local centers in the global coreset may select nearly identical points at different nodes if the local datasets are similar, which is non-optimal and inefficient.

An existing algorithm: The state-of-the-art solution to the distributed k -clustering problem is based on an algorithm called *Communication-aware Distributed Coreset Construction (CDCC)* [7]. In our context, this solution works as follows:

- 1) the server allocates a given sample size t among the nodes, such that the sample size t_j allocated to node v_j is proportional to the local k -clustering cost $c(P_j, B_j)$ reported by v_j ;
- 2) each node v_j generates and reports a local coreset D_j , consisting of the local centers B_j and t_j points sampled i.i.d. from P_j , where each $p \in P_j$ has a sampling probability proportional to the cost of clustering p to the nearest center in B_j ;
- 3) the server computes a set of k -clustering centers Q_D from the coreset $D = \bigcup_{j=1}^n D_j$.

It is shown in [7] that if $t = O(\frac{1}{\epsilon^2}(kd + \log \frac{1}{\delta}))$ for k -median and $t = O(\frac{1}{\epsilon^4}(kd + \log \frac{1}{\delta}) + nk \log \frac{nk}{\delta})$ for k -means, then with probability at least $1 - \delta$, D is an ϵ -coreset for P w.r.t. the cost function of k -median/means. According to Definition II.1, this implies that if Q_P is the set of optimal k -clustering centers for P , then $c(P, Q_D)/c(P, Q_P) \leq (1 + \epsilon)/(1 - \epsilon)$.

Adaptation for coreset construction: First, we skip step (3) (i.e., computation of Q_D) and directly use $D = \bigcup_{j=1}^n D_j$ as the coreset. This is because the coreset of a coreset cannot have a better quality than the original coreset [40].

Moreover, in CDCC, the number of local centers k is a given parameter as it is only designed to support k -clustering. Since our goal is to support a variety of machine learning problems, the number of local centers k_j at each v_j becomes a design parameter that can vary across nodes. Given a global coreset size N , we will show that the approximation error of the constructed coreset depends on $(k_j)_{j=1}^n$ through $\frac{1}{\sqrt{N - \sum_{j=1}^n k_j}} \sum_{j=1}^n \text{approx}(P_j, k_j)$ (see Theorem IV.1). Thus, we set $(k_j)_{j=1}^n$ to minimize this error, and obtain the remaining $t = N - \sum_{j=1}^n k_j$ points by sampling.

Combining these ideas yields a distributed coreset construction algorithm called *Distributed Robust Coreset Construction (DRCC)* (Algorithm 2). The algorithm works in three steps: (1) each node reports its local k -clustering cost for a range of k (lines 2-3), (2) the server uses the reported costs to configure the number of local centers k_j and the number of random samples t_j at each node v_j (lines 5-7), and (3) each node independently constructs a local coreset using a combination of samples and local centers (lines 9-12). DRCC generalizes CDCC in that: (i) it allows the input dataset to be weighted (w_p : weight of input point p ; u_q : weight of coreset point q); (ii) it allows the number of local centers to be different for different nodes. In the special case of $k_j \equiv k$ for all $j \in [n]$ and $w_p \equiv 1$ for all $p \in P$, DRCC is reduced to CDCC.

Communication overhead: DRCC has a communication overhead of $O(Kn)$, measured by the total number of scalars reported by the nodes besides the coreset itself. In practice, K should be a small constant to allow efficient computation of local k -clustering for $k \in [K]$. This overhead is much smaller than the $O((n-1)Nd)$ overhead of the merge-and-reduce

Algorithm 2: Distributed Robust Coreset Construction
 $((P_j)_{j=1}^n, N, K)$

input : A distributed dataset $(P_j)_{j=1}^n$, global coreset size N , maximum number of local centers K

output: A coreset $D = \bigcup_{j=1}^n (S_j \cup B_j^{k_j})$ for $P = \bigcup_{j=1}^n P_j$

- 1 **each** v_j ($j \in [n]$):
 - 2 compute local approximate k -clustering centers B_j^k
 on P_j for $k = 1, \dots, K$;
 - 3 report $(c(P_j, B_j^k))_{k=1}^K$ to the server;
 - 4 **the server:**
 - 5 find $(k_j)_{j=1}^n$ that minimizes
 $\frac{1}{\sqrt{N - \sum_{j=1}^n k_j}} \sum_{j=1}^n c(P_j, B_j^{k_j})$ s.t. $k_j \in [K]$ and
 $\sum_{j=1}^n k_j \leq N$;
 - 6 randomly allocate $t = N - \sum_{j=1}^n k_j$ points i.i.d.
 among v_1, \dots, v_n , where each point belongs to v_j
 with probability $\frac{c(P_j, B_j^{k_j})}{\sum_{j=1}^n c(P_j, B_j^{k_j})}$;
 - 7 communicate $(k_j, t_j, \frac{C}{t})$ to each v_j ($j \in [n]$),
 where t_j is the number of points allocated to v_j
 and $C = \sum_{l=1}^n c(P_l, B_l^{k_l})$;
 - 8 **each** v_j ($j \in [n]$):
 - 9 sample a set S_j of t_j points i.i.d. from P_j , where
 each sample equals $p \in P_j$ with probability
 $\frac{m_p}{c(P_j, B_j^{k_j})}$ for $m_p = c(\{p\}, B_j^{k_j})$;
 - 10 set the weight of each $q \in S_j$ to $u_q = \frac{C w_q}{t m_q}$;
 - 11 set the weight of each $b \in B_j^{k_j}$ to
 $u_b = \sum_{p \in P_b} w_p - \sum_{q \in P_b \cap S_j} u_q$, where P_b is the
 set of points in P_j whose closest center in $B_j^{k_j}$ is
 b ;
 - 12 report each point $q \in S_j \cup B_j^{k_j}$ and its weight u_q
 to the server;
-

approach² in [40].

Quality of coreset: Regarding the quality of the coreset, we have proved the following result in Appendix A. Given a general per-point cost function $\text{cost}(p, x)$, define $f_x(p) := w_p(\text{cost}(p, x) - \text{cost}(b_p, x) + \rho \text{dist}(p, b_p))$, where b_p is the center in $B_j^{k_j}$ closest to $p \in P_j$. Let $\dim(F, P)$ denote the dimension of the function space $F := \{f_x(p) : x \in \mathcal{X}\}$ [7].

Theorem IV.1. *If $\text{cost}(p, x)$ is ρ -Lipschitz-continuous in p for any $x \in \mathcal{X}$, then $\exists t = O(\frac{1}{\epsilon^2}(\dim(F, P) + \log \frac{1}{\delta}))$ such that with probability at least $1 - \delta$, the coreset D constructed by DRCC based on local k -median clustering, which contains k_j*

²Specifically, the merge-and-reduce approach works by applying a centralized coreset construction algorithm (e.g., Algorithm 1) repeatedly to combine local coresets into a global coreset. Given n local coresets computed by each of the n nodes, each containing N points in \mathbb{R}^d , the centralized algorithm needs to be applied $n - 1$ times, each time requiring a local coreset to be transmitted to the location of another local coreset. This results in a total communication overhead of $(n - 1)Nd$.

local centers from v_j ($j \in [n]$) and t random samples, satisfies

$$\left| \sum_{p \in P} w_p \text{cost}(p, x) - \sum_{q \in D} u_q \text{cost}(q, x) \right| \leq 2\epsilon \rho \sum_{j=1}^n c(P_j, B_j^{k_j}) \quad (13)$$

for all $x \in \mathcal{X}$.

Here, the parameter $\dim(F, P)$ is a property of the machine learning problem under consideration, which intuitively measures the degree of freedom in the solution $x \in \mathcal{X}$, e.g., $\dim(F, P) = O(kd)$ for k -means/median in d -dimensional space [23]. See Appendix C for more discussions.

Due to the relationship between t and ϵ given in Theorem IV.1, the bound on the right-hand side of (13) depends on parameters t and $(k_j)_{j=1}^n$ through $\frac{1}{\sqrt{t}} \sum_{j=1}^n c(P_j, B_j^{k_j})$. Specifically, given a total coreset size N , the right-hand side of (13) is

$$O \left(\frac{\rho \sqrt{\dim(F, P) + \log \frac{1}{\delta}}}{\sqrt{N - \sum_{j=1}^n k_j}} \cdot \sum_{j=1}^n c(P_j, B_j^{k_j}) \right). \quad (14)$$

This error bound tells us that the approximation error decreases with the coreset size N at roughly $O(1/\sqrt{N})$. The error, however, may not be monotone with the numbers of local centers k_j 's, as increasing their values decreases both $\sqrt{N - \sum_{j=1}^n k_j}$ and $\sum_{j=1}^n c(P_j, B_j^{k_j})$. Thus, we select $(k_j)_{j=1}^n$ to minimize the error bound in line 5 of Algorithm 2. As the server needs to know $(c(P_j, B_j^k))_{k=1}^K$ ($\forall j \in [n]$) to solve this minimization over $k_j \in [K]$, the choice of the parameter K faces a tradeoff: a larger K yields a larger solution space and possibly a better configuration of $(k_j)_{j=1}^n$ to minimize the approximation error, but incurs a higher communication (and computation) overhead at the nodes. The optimal K will depend on the desirable tradeoff and the specific dataset.

Remark: The performance bound in Theorem IV.1 is on the absolute error, instead of the relative error as guaranteed by an ϵ -coreset. Nevertheless, if $\exists \beta > 0$ and $(k_j)_{j=1}^n$ such that $\text{cost}(P, x) = \sum_{p \in P} w_p \text{cost}(p, x) \geq \beta \sum_{j=1}^n c(P_j, B_j^{k_j})$ for all $x \in \mathcal{X}$, then (13) implies that D is an ϵ -coreset for P w.r.t. $\text{cost}(P, x)$ with probability at least $1 - \delta$ if $t = O(\frac{\rho^2}{\epsilon^2}(\dim(F, P) + \log \frac{1}{\delta}))$, i.e., the total coreset size $N = O(\frac{\rho^2}{\epsilon^2}(\dim(F, P) + \log \frac{1}{\delta}) + \sum_{j=1}^n k_j)$. In the special case where $\text{cost}(P, x)$ is the k -median clustering cost and $k_j \equiv k$, we have $\rho = 1$ (Appendix B) and $\dim(F, P) = O(kd)$ [23], and thus the size of an ϵ -coreset is $O(\frac{1}{\epsilon^2}(kd + \log \frac{1}{\delta}) + kn)$, which generalizes the result in [7] to weighted datasets.

V. PERFORMANCE EVALUATION

We evaluate the proposed coreset construction algorithms and their benchmarks on a variety of machine learning problems and datasets, and compare the cost of each model learned on a coreset with the cost of the model learned on the original dataset. We first perform evaluations in a centralized setting to compare different approaches to construct coresets, and then evaluate different ways of applying the most promising approach in a distributed setting.

Table I
PARAMETERS OF DATASETS

dataset	size ($ P $)	dimension (d)	#distinct labels (L)
Fisher's iris	150	5	3
Facebook	500	19	4
Pendigits	7494	17	10
MNIST	70000	401	10
HAR	10299	562	6

Coreset construction algorithms: In the centralized setting, we evaluate RCC based on k -median clustering ('RCC-kmedian') and RCC based on k -means clustering ('RCC-kmeans'), together with benchmarks including the algorithm in [24] ('farthest point'), the framework in [23] instantiated for k -means ('nonuniform sampling'), and uniform sampling ('uniform sampling'). We note that the algorithm in [9] ('decomposition' in Figure 2) is essentially RCC based on k -means clustering, hence omitted.

In the distributed setting, we take the best-performing algorithm in the centralized setting ('RCC-kmeans') and evaluate its distributed extensions – including CDCC [7] and DRCC.

Datasets: We use five real datasets: (1) Fisher's iris data [46], which is a 5-dimensional dataset consisting of measurements of 150 iris specimens and their species, (2) Facebook metrics [47], which is a 19-dimensional dataset consisting of features of 500 posts published in a popular Facebook page, (3) Pendigits data [48], which is a 17-dimensional dataset consisting of feature vectors of 7494 handwritten digits, (4) MNIST data [49], which consists of 60,000 images of handwritten digits for training plus 10,000 images for testing, each trimmed to 20×20 pixels, and (5) Human Activity Recognition (HAR) using Smartphones data [50], which contains 10,299 samples of smartphone sensor readings during 6 different activities, each sample containing 561 readings.

We normalize each numerical dimension to $[0, 1]$. We map labels to numbers such that the distance between two points with different labels is no smaller than the distance between points with the same label. Given a d -dimensional dataset (including labels) with L types of labels, we map type- l label to $(l-1)\tau$ ($l \in [L]$) for $\tau = \lceil \sqrt{d-1} \rceil$. See Table I for a summary. In testing SVM, we map one label to '1' and the rest to '-1'. Each data point has a unit weight.

To generate distributed datasets, we use three schemes: (i) *uniform*, where the points are uniformly distributed across n nodes, (ii) *specialized*, where each node is associated with one label and contains all the data points with this label, and (iii) *hybrid*, where the first n_0 nodes are "specialized" as in (ii), and the remaining data are randomly partitioned among the remaining nodes.

Machine learning problems: We evaluate three unsupervised learning problems—MEB, k -means, and PCA, and two supervised learning problem—SVM and Neural Network (NN). In our NN experiment, we define a three-layer network, whose hidden layer has 100 neurons. As an NN of this size

³The model x denotes the center of enclosing ball for MEB and the set of centers for k -means. For PCA, $x = WW^T$, where W is a $d \times l$ matrix consisting of the first l ($l < d$) principle components as columns. For SVM, $x_{1:d-1} \in \mathbb{R}^{d-1}$ is the coefficient vector and $x_d \in \mathbb{R}$ is the offset.

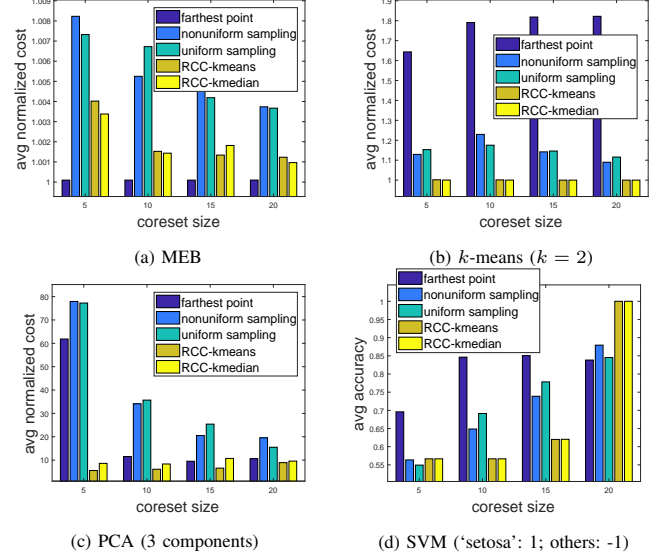


Figure 3. Evaluation on Fisher's iris dataset with varying coreset size (label: 'species').

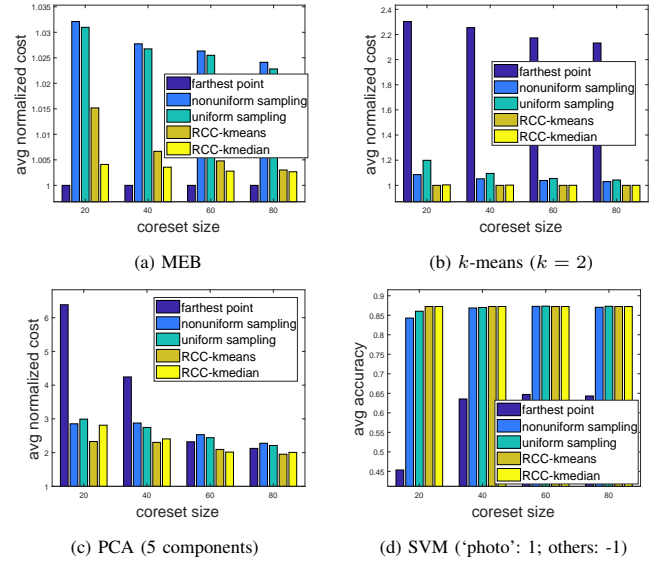



Figure 4. Evaluation on Facebook metrics dataset with varying coreset size (label: 'type').



is typically used on datasets with at least 100 features, we only evaluate NN on MNIST and HAR, and replace it by SVM for the other datasets. Table II gives their cost functions, where for a data point $p \in \mathbb{R}^d$, $p_{1:d-1} \in \mathbb{R}^{d-1}$ denotes the numerical portion and $p_d \in \mathbb{R}$ denotes the label. The meaning of the model parameter x is problem-specific, as explained in the footnote. We also provide (upper bounds of) the Lipschitz constant ρ except for NN, since it is NP-hard to evaluate ρ for even a two-layer network [45]; see Appendix B for analysis. Here l is the number of principle components computed by PCA, and Δ is the diameter of the sample space. In our experiments, $\Delta = \sqrt{(d-1)(L^2 - 2L + 2)}$, which is 4.5 for Fisher's iris, 13.4 for Facebook, 36.2 for Pendigits, 181.1 for MNIST, and 120.8 for HAR. While SVM and NN do not have a meaningful ρ , we still include them to stress-test our algorithm.

Table II
MACHINE LEARNING COST FUNCTIONS


problem	overall cost function ³	ρ
MEB	$\max_{p \in P} \text{dist}(x, p)$	1
k -means	$\sum_{p \in P} w_p \cdot \min_{q_i \in x} \text{dist}(q_i, p)^2$	2Δ
PCA	$\sum_{p \in P} w_p \cdot \text{dist}(p, xp)^2$	$2\Delta(l+1)$
SVM	$\sum_{p \in P} w_p \max(0, 1 - p_d(p_{1:d-1}^T x_{1:d-1} + x_d))$	∞
Neural Net	$\sum_{p \in P} (-p_d) \cdot \log(o_p)$, where o_p is the output for input $p_{1:d-1}$	NP-hard [45]




(a) MEB

(b) k -means ($k = 2$)


(c) PCA (11 components)



(d) SVM ('0': 1; others: -1)




(e) PCA (7 components)





(f) NN


Figure 5. Evaluation on Pendigits with varying coreset size (label: 'digit').



(a) MEB

(b) k -means ($k = 2$)

(c) PCA (300 components)




(d) NN



Figure 6. Evaluation on MNIST with varying coreset size (label: 'labels').

Performance metrics: For the unsupervised learning problems (MEB, k -means, and PCA), we evaluate the performance by the normalized cost as explained in Section III-A. For the supervised learning problems (SVM and NN), we evaluate the performance by the accuracy in predicting the labels of testing data. MNIST and HAR datasets are already divided into training set and testing set. For other datasets, we use the first 80% of data for training, and the rest for testing.


Figure 7. Evaluation on HAR with varying coreset size (label: 'labels').



(a) MEB

(b) k -means ($k = 2$)

(c) PCA (3 components)



(d) SVM ('setosa': 1; others: -1)

Figure 8. Detailed evaluation on Fisher's iris dataset (label: 'species', coreset size: 20).

Results in centralized setting: Figures 3–7 show the performances achieved at a variety of coreset sizes, averaged over 100 Monte Carlo runs. Better performance is indicated by lower cost for an unsupervised learning problem or higher accuracy for a supervised learning problem. Note that even the largest coresets generated in these experiments are much smaller (by 84–99.3%) than the original dataset, implying

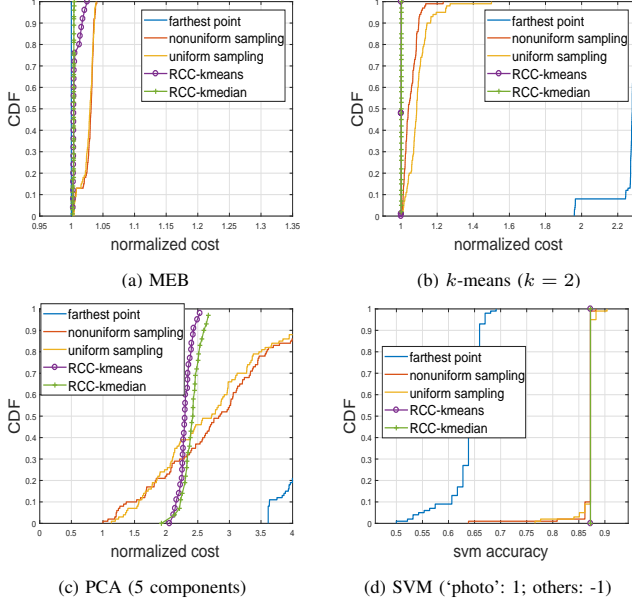


Figure 9. Detailed evaluation on Facebook metrics dataset (label: 'type', coreset size: 40).

Table III
AVERAGE RUNNING TIME (SEC) ('FP': FARTEST POINT, 'NS': NONUNIFORM SAMPLING, 'US': UNIFORM SAMPLING, 'RS': RCC-KMEANS, 'RN': RCC-KMEDIAN)

algorithm	Fisher	Facebook	Pendigits	MNIST	HAR
FP	1.62	3.00	2.53	21.69	25.92
NS	0.019	0.027	0.095	7.42	0.69
US	2.10e-04	4.60e-04	3.80e-04	0.01	0.0013
RS	0.0083	0.011	0.042	18.76	1.46
RN	0.028	0.30	0.40	100.64	12.39

significant reduction in the communication cost by reporting a coreset instead of the raw data.

We see that the proposed algorithms ('RCC-kmeans' and 'RCC-kmedian') perform either the best or comparably to the best across all the datasets and all the machine learning problems. The farthest point algorithm in [24], designed for MEB, can perform very poorly for other machine learning problems. The sampling-based algorithms ('nonuniform sampling' [23] and 'uniform sampling') perform relatively poorly for MEB and PCA. Generally, we see that the advantages of RCC algorithms are more significant at small coreset sizes. One exception is the SVM accuracy for Fisher's iris (Figure 3 (d)), where points on the peripheral of the dataset (which are likely to be chosen by the farthest point algorithm) happen to have different labels and induce a rough partition between the points labeled '1' and those labeled '-1', causing better performance for 'farthest point' at very small coreset sizes.

Besides the average normalized costs, we also evaluated the CDFs of the results, shown in Figures 8–12. The results show similar comparisons as observed before. Moreover, we see that the proposed algorithms ('RCC-kmeans' and 'RCC-kmedian') also have significantly less performance variation than the benchmarks, especially the sampling-based algorithms ('nonuniform sampling' and 'uniform sampling'). This means

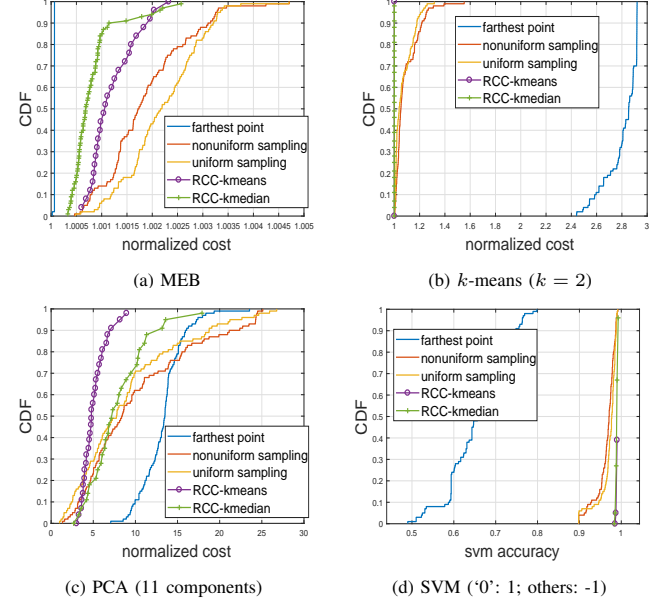


Figure 10. Detailed evaluation on Pendigits dataset (label: 'digit', coreset size: 40).

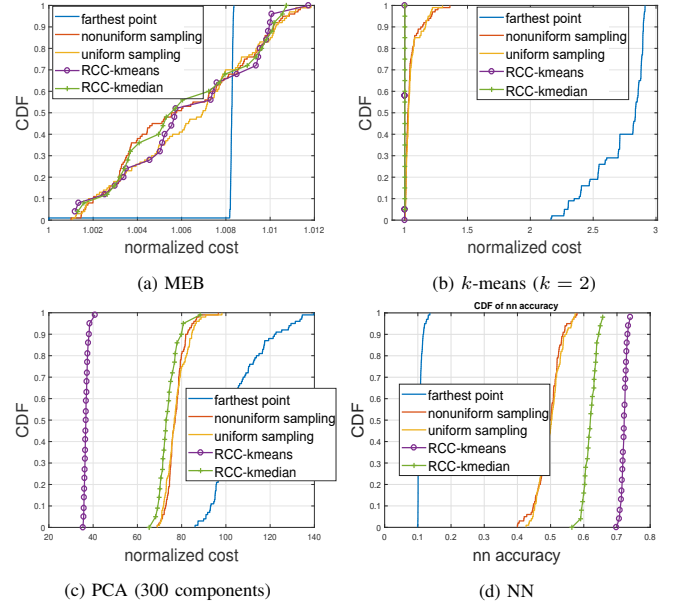


Figure 11. Detailed evaluation on MNIST dataset (label: 'labels', coreset size: 50).

that the quality of the coresets constructed by the proposed algorithms is more reliable, which is a desirable property.

Between the proposed algorithms, 'RCC-kmeans' sometimes outperforms 'RCC-kmedian', e.g., Figure 6 (c–d). Moreover, we note that 'RCC-kmeans' can be an order of magnitude faster than 'RCC-kmedian', as shown in Table III. Note that our primary goal in constructing a robust coreset is to reduce the communication cost in scenarios like Figure 1 while supporting diverse machine learning problems, instead of speeding up the coreset construction at the data source. This result shows that such robustness may come with certain penalty in running time. Nevertheless, the running time of 'RCC-kmeans' is comparable to the benchmarks.

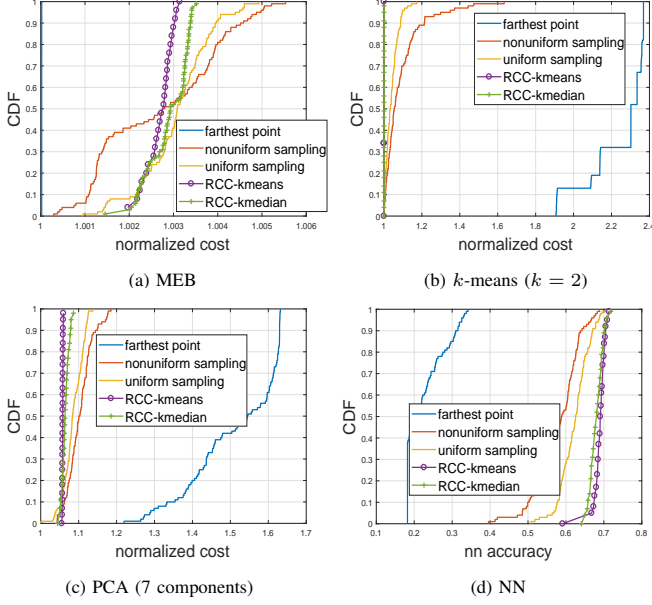


Figure 12. Detailed evaluation on HAR dataset (label: ‘labels’, coreset size: 50).

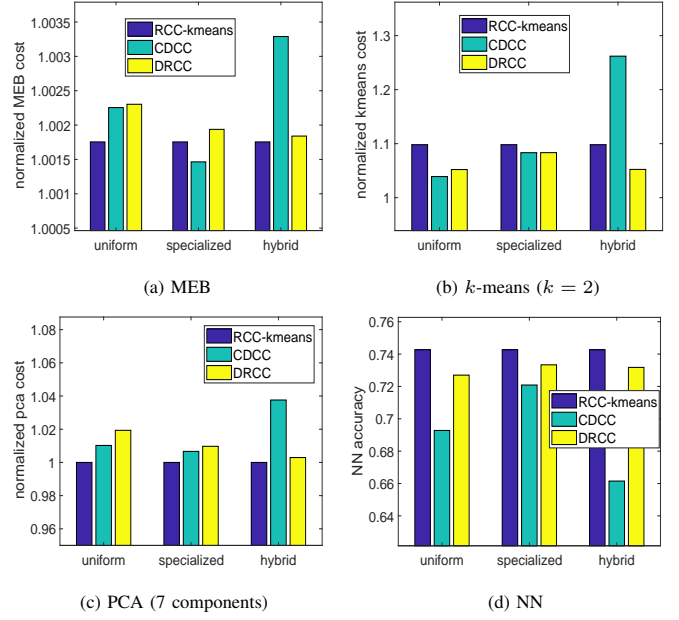


Figure 14. Evaluation on HAR in distributed setting (label: ‘activity’, coreset size: 200, $K = 10$).




Figure 13. Evaluation on MNIST in distributed setting (label: ‘labels’, coreset size: 400, $K = 10$).

Results in distributed setting: We use MNIST and HAR to generate distributed datasets at $n = L$ nodes according to the three aforementioned schemes (‘uniform’, ‘specialized’, ‘hybrid’), where node j ($j \in [n]$) is associated with label ‘ $j - 1$ ’, and $n_0 = \lfloor n/2 \rfloor$. We parameterize CDCC with $k = 2$ according to the evaluated k -means problem. For DRCC (Algorithm 2), we solve line 5 by a greedy heuristic. As a bound, we also show the performance of RCC based on k -means clustering (‘RCC-kmeans’), which is the best-performing algorithm in the centralized setting. We generate 5 distributed datasets using each scheme and repeat coreset construction for 5 times on each dataset. Figures 13–14 show the average results under fixed coreset size N and parameter

Table IV
AVERAGE RUNNING TIME (SEC)

algorithm	MNIST	HAR
CDCC	13.84	1.55
DRCC	31.66	2.42

K (defined in Algorithm 2); see Appendix D for additional evaluation results when we vary these parameters.

As CDCC blindly configures k coreset points to be local centers at each node, it suffers when the local datasets are highly heterogeneous (under the ‘hybrid’ scheme). By automatically tuning the numbers of local centers according to the local clustering costs, DRCC is able to customize its configuration to the dataset at hand and thus achieves a more robust performance, especially under the ‘hybrid’ scheme. We note that as CDCC was designed to support k -means, its performance for k -means can be slightly better than DRCC (e.g., Figure 13–14 (b) under the ‘uniform’ scheme), but the difference is very small ($< 3\%$). Meanwhile, the robustness of DRCC comes with a small cost in communication overhead, where each node needs to report K scalars (here $K = 10$) instead of one scalar as in CDCC. However, this difference is negligible compared to the cost of reporting the coreset itself (e.g., 400×401 scalars for MNIST and 200×562 scalars for HAR). Both algorithms generate coresets that are much smaller (by 97.2–99.3%) than the original dataset, significantly reducing the communication cost. Note that ‘RCC-kmeans’ is a centralized algorithm that is only used as a benchmark here.

Additionally, Table IV shows the average running times of CDCC and DRCC on each dataset. Similar to the centralized setting, the better robustness of DRCC comes with certain penalty in running time due to the need to compute multiple instances of k -clustering ($k = 1, \dots, K$) on local datasets.

VI. CONCLUSION

We show, both theoretically and empirically, that the k -clustering centers form a coreset that provides a guaranteed approximation for a broad set of machine learning problems with sufficiently continuous cost functions. As k -clustering (including k -means/median) is one of the most well-studied machine learning problems, this result allows us to leverage existing k -clustering algorithms for robust coreset construction. In particular, we have adapted an existing distributed k -clustering algorithm to construct a robust coreset over distributed data with a very low communication overhead. Our extensive evaluations on real datasets verify the superior performance of the proposed algorithms in supporting diverse machine learning problems.

APPENDIX A: PROOFS

A.1. Proof of Lemma III.1:

Proof. By definition, $\text{opt}(P, k) = \sum_{i=1}^k \text{opt}(P_i, 1)$. By letting $X_i := \{X_{i,1}, X_{i,2}\}$ be the centers of the optimal 2-clustering of P_i , we have

$$\begin{aligned} \sum_{i=1}^k \text{opt}(P_i, 2) &= \sum_{i=1}^k \sum_{p \in P_i} w_p \left(\min_{x \in X_i} \text{dist}(p, x) \right)^z \\ &\geq \sum_{i=1}^k \sum_{p \in P_i} w_p \left(\min_{x \in \bigcup_{i=1}^k X_i} \text{dist}(p, x) \right)^z \geq \text{opt}(P, 2k). \end{aligned} \quad (15)$$

Thus we have

$$\sum_{i=1}^k (\text{opt}(P_i, 1) - \text{opt}(P_i, 2)) \leq \text{opt}(P, k) - \text{opt}(P, 2k) \leq \epsilon'. \quad (16)$$

Since $\text{opt}(P_i, 1) - \text{opt}(P_i, 2) \geq 0$, $\text{opt}(P_i, 1) - \text{opt}(P_i, 2) \leq \epsilon'$, $\forall i \in [k]$. \square

A.2. Proof of Lemma III.2:

Proof. By definition of k -clustering, we have

$$\text{opt}(P_i, 1) = \sum_{p \in P_i} w_p \text{dist}(p, \mu(P_i))^z, \quad (17)$$

$$\text{opt}(P_i, 2) \leq \sum_{p \in P_i^1} w_p \text{dist}(p, \mu(P_i))^z + \sum_{p \in P_i^2} w_p \text{dist}(p, p_0)^z \quad (18)$$

for any $p_0 \in P_i$, where P_i^1 is the subset of points in P_i that are closer to $\mu(P_i)$ than p_0 (ties broken arbitrarily) and $P_i^2 := P_i \setminus P_i^1$. Then subtracting (17) by (18), we have

$$\begin{aligned} \text{opt}(P_i, 1) - \text{opt}(P_i, 2) &\geq \sum_{p \in P_i^2} w_p (\text{dist}(p, \mu(P_i))^z - \text{dist}(p, p_0)^z). \end{aligned} \quad (19)$$

As $\text{dist}(p, \mu(P_i))^z - \text{dist}(p, p_0)^z \geq 0$ for all $p \in P_i^2$,

$$\begin{aligned} &w_p (\text{dist}(p, \mu(P_i))^z - \text{dist}(p, p_0)^z) \\ &\leq \text{opt}(P_i, 1) - \text{opt}(P_i, 2) \\ &\leq \epsilon', \quad \forall p \in P_i^2. \end{aligned} \quad (20)$$

In particular, for $p = p_0$, $w_{p_0} \text{dist}(p_0, \mu(P_i))^z \leq \epsilon'$ and thus $\text{dist}(p_0, \mu(P_i)) \leq (\frac{\epsilon'}{w_{\min}})^{\frac{1}{z}}$. The proof completes by noting that this holds for any $p_0 \in P_i$. \square

A.3. Proof of Lemma III.3:

Proof. Since $\forall x \in \mathcal{X}$ and $p \in P_i$,

$$(1 - \epsilon) \text{cost}(p, x) \leq \text{cost}(\mu(P_i), x) \leq (1 + \epsilon) \text{cost}(p, x), \quad (21)$$

multiplying this inequality by w_p and then summing over $p \in P_i$, we have

$$\begin{aligned} (1 - \epsilon) \sum_{p \in P_i} w_p \text{cost}(p, x) &\leq \text{cost}(\mu(P_i), x) \sum_{p \in P_i} w_p \\ &\leq (1 + \epsilon) \sum_{p \in P_i} w_p \text{cost}(p, x), \quad \forall x \in \mathcal{X}. \end{aligned} \quad (22)$$

Summing up (22) over all $i \in [k]$, we have

$$(1 - \epsilon) \text{cost}(P, x) \leq \text{cost}(S, x) \leq (1 + \epsilon) \text{cost}(P, x), \quad \forall x \in \mathcal{X}.$$

Therefore, S is an ϵ -coreset for P w.r.t. $\text{cost}(P, x)$. \square

A.4. Proof of Lemma III.4

Proof. Taking maximum over $p \in P_i$ for (5), we have that $\forall i \in [k]$ and $x \in \mathcal{X}$,

$$\begin{aligned} (1 - \epsilon) \max_{p \in P_i} \text{cost}(p, x) &\leq \text{cost}(\mu(P_i), x) \\ &\leq (1 + \epsilon) \max_{p \in P_i} \text{cost}(p, x). \end{aligned} \quad (23)$$

Taking maximum over $i \in [k]$ for (23): $\forall x \in \mathcal{X}$,

$$\begin{aligned} (1 - \epsilon) \max_{p \in P} \text{cost}(p, x) &\leq \max_{i \in [k]} \text{cost}(\mu(P_i), x) \\ &\leq (1 + \epsilon) \max_{p \in P} \text{cost}(p, x). \end{aligned}$$

That is, S is an ϵ -coreset for P w.r.t. $\text{cost}(P, x)$. \square

A.5. Proof of Lemma III.5:

Proof. By definition and Assumption 2, we have

$$\text{approx}(P, k) = \sum_{i=1}^k \text{approx}(\tilde{P}_i, 1) \quad (24)$$

$$\sum_{i=1}^k \text{approx}(\tilde{P}_i, 2) \geq \text{approx}(P, 2k) \quad (25)$$

Thus, we have

$$\begin{aligned} &\sum_{i=1}^k (\text{approx}(\tilde{P}_i, 1) - \text{approx}(\tilde{P}_i, 2)) \\ &\leq \text{approx}(P, k) - \text{approx}(P, 2k) \leq \epsilon'. \end{aligned} \quad (26)$$

Let $\{\tilde{P}_i^1, \tilde{P}_i^2\}$ be the partition of \tilde{P}_i generated by the algorithm for 2-clustering of \tilde{P}_i . Under Assumption 1,

$$\begin{aligned} &\text{approx}(\tilde{P}_i, 2) \\ &= \sum_{p \in \tilde{P}_i^1} w_p \text{dist}(p, \mu(\tilde{P}_i^1))^z + \sum_{p \in \tilde{P}_i^2} w_p \text{dist}(p, \mu(\tilde{P}_i^2))^z \end{aligned}$$

$$\begin{aligned} &\leq \sum_{p \in \tilde{P}_i^1} w_p \text{dist}(p, \mu(\tilde{P}_i))^z + \sum_{p \in \tilde{P}_i^2} w_p \text{dist}(p, \mu(\tilde{P}_i))^z \\ &= \text{approx}(\tilde{P}_i, 1). \end{aligned} \quad (27)$$

Combining (26, 27) yields

$$\text{approx}(\tilde{P}_i, 1) - \text{approx}(\tilde{P}_i, 2) \leq \epsilon', \forall i \in [k]. \quad (28)$$

□

A.6. Proof of Lemma III.6:

Proof. First, by Assumption 1, $\text{approx}(\tilde{P}_i, 1) = \sum_{p \in \tilde{P}_i} w_p \cdot \text{dist}(p, \mu(\tilde{P}_i))^z$. Moreover, by Assumption 3,

$$\begin{aligned} \text{approx}(\tilde{P}_i, 2) &\leq \sum_{p \in \tilde{P}_i} w_p \left(\min_{x \in \{\mu(\tilde{P}_i), p^*\}} \text{dist}(p, x) \right)^z \\ &\leq \sum_{p \in \tilde{P}_i \setminus \{p^*\}} w_p \text{dist}(p, \mu(\tilde{P}_i))^z, \end{aligned} \quad (29)$$

where $p^* := \arg \max_{p \in \tilde{P}_i} w_p \text{dist}(p, \mu(\tilde{P}_i))^z$. Thus, we have

$$\begin{aligned} \epsilon' &\geq \text{approx}(\tilde{P}_i, 1) - \text{approx}(\tilde{P}_i, 2) \\ &\geq w_{p^*} \text{dist}(p^*, \mu(\tilde{P}_i))^z \geq w_p \text{dist}(p, \mu(\tilde{P}_i))^z, \forall p \in \tilde{P}_i. \end{aligned} \quad (30)$$

Therefore, $\text{dist}(p, \mu(\tilde{P}_i)) \leq (\frac{\epsilon'}{w_{\min}})^{\frac{1}{z}}, \forall p \in \tilde{P}_i$. □

A.7. Proof of Theorem IV.1:

Proof. The proof is based on a sampling lemma from [7]:

Lemma A.1 ([7]). *Let F be a set of nonnegative functions $f : P \rightarrow \mathbb{R}_{\geq 0}$. Let S be a set of i.i.d. samples from P , where each sample equals $p \in P$ with probability $\frac{m_p}{\sum_{z \in P} m_z}$.*

Let $u'_q = \frac{\sum_{p \in P} m_p}{m_q |S|}$ for $q \in P$. If for a sufficiently large constant α , $|S| \geq \frac{\alpha}{\epsilon^2} (\dim(F, P) + \log \frac{1}{\delta})$, then with probability at least $1 - \delta$, $\forall f \in F : |\sum_{p \in P} f(p) - \sum_{q \in S} u'_q f(q)| \leq \epsilon (\sum_{p \in P} m_p) (\max_{p \in P} \frac{f(p)}{m_p})$.

In our case, $S = \bigcup_{j=1}^n S_j$, and $B = \bigcup_{j=1}^n B_j^{k_j}$. Define $f_x(p) := w_p (\text{cost}(p, x) - \text{cost}(b_p, x) + \rho \text{dist}(p, b_p))$, where b_p is the center in $B_j^{k_j}$ closest to $p \in P_j$. By the ρ -Lipschitz-continuity of $\text{cost}(p, x)$, $f_x(p) \geq 0$ and $f_x(p) \leq 2\rho w_p \text{dist}(p, b_p)$. By lines 6 and 9 in Algorithm 2, k -median-based DRCC generates S via i.i.d. sampling from P with probabilities proportional to $m_p = w_p \text{dist}(p, b_p)$. Therefore, by Lemma A.1, there exists $|S| = O(\frac{1}{\epsilon^2} (\dim(F, P) + \log \frac{1}{\delta}))$ such that with probability at least $1 - \delta$, $\forall x \in \mathcal{X}$:

$$\Delta := \left| \sum_{p \in P} f_x(p) - \sum_{q \in S} u'_q f_x(q) \right| \leq \epsilon \left(\sum_{p \in P} m_p \right) \left(\max_{p \in P} \frac{f_x(p)}{m_p} \right). \quad (31)$$

The righthand side of (31) equals the righthand side of (13) as $\sum_{p \in P} m_p = \sum_{j=1}^n c(P_j, B_j^{k_j})$ and $\frac{f_x(p)}{m_p} \leq 2\rho$.

We will show that the lefthand side of (31) also equals the lefthand side of (13). Specifically,

$$\sum_{p \in P} f_x(p) = \sum_{p \in P} w_p \text{cost}(p, x) - \sum_{b \in B} \text{cost}(b, x) \sum_{p \in P_b} w_p$$

$$+ \rho \sum_{p \in P} m_p, \quad (32)$$

and since $u'_q w_q = u_q$ for each $q \in S$ (line 10 in Algorithm 2),

$$\begin{aligned} \sum_{q \in S} u'_q f_x(q) &= \sum_{q \in S} u_q \text{cost}(q, x) \\ &- \sum_{b \in B} \text{cost}(b, x) \sum_{q \in P_b \cap S} u_q + \rho \sum_{q \in S} u_q \text{dist}(q, b_q). \end{aligned} \quad (33)$$

Since $\sum_{q \in S} u_q \text{dist}(q, b_q) = \sum_{p \in P} m_p$,

$$\begin{aligned} \Delta &= \left| \sum_{p \in P} w_p \text{cost}(p, x) - \sum_{q \in S} u_q \text{cost}(q, x) \right. \\ &\quad \left. - \sum_{b \in B} \text{cost}(b, x) \left(\sum_{p \in P_b} w_p - \sum_{q \in P_b \cap S} u_q \right) \right| \\ &= \left| \sum_{p \in P} w_p \text{cost}(p, x) - \sum_{q \in S \cup B} u_q \text{cost}(q, x) \right|, \end{aligned} \quad (34)$$

as $u_b = \sum_{p \in P_b} w_p - \sum_{q \in P_b \cap S} u_q$ (line 11 in Algorithm 2). □

APPENDIX B: ANALYSIS OF LIPSCHITZ CONSTANT

• For MEB, $\text{cost}(p, x) = \text{dist}(p, x)$, where $x \in \mathbb{R}^d$ denotes the center of the enclosing ball. For any data points $p, p' \in \mathbb{R}^d$, by triangle inequality, we have:

$$|\text{dist}(p, x) - \text{dist}(p', x)| \leq \text{dist}(p, p'). \quad (35)$$

Hence, its cost function is 1-Lipschitz-continuous ($\rho = 1$).

• For k -median, $\text{cost}(p, x) = \min_{q \in x} \text{dist}(p, q)$, where $x \subset \mathbb{R}^d$ denotes the set of k centers. For any data points $p, p' \in \mathbb{R}^d$, without loss of generality (WLOG), suppose $\min_{q \in x} \text{dist}(p, q) \geq \min_{q \in x} \text{dist}(p', q) = \text{dist}(p', q')$ for some $q' \in x$. Then

$$\begin{aligned} &|\min_{q \in x} \text{dist}(p, q) - \min_{q \in x} \text{dist}(p', q)| \\ &= \min_{q \in x} \text{dist}(p, q) - \text{dist}(p', q') \\ &\leq \text{dist}(p, q') - \text{dist}(p', q') \leq \text{dist}(p, p'). \end{aligned} \quad (36)$$

Hence, $\rho = 1$ for k -median.

• For k -means, $\text{cost}(p, x) = \min_{q \in x} \text{dist}(p, q)^2$, where x denotes the set of k centers. Similar to the above, for any data points $p, p' \in \mathbb{R}^d$, suppose WLOG that $\min_{q \in x} \text{dist}(p, q) \geq \min_{q \in x} \text{dist}(p', q) = \text{dist}(p', q')$ for some $q' \in x$. Then

$$\begin{aligned} &|\min_{q \in x} \text{dist}(p, q)^2 - \min_{q \in x} \text{dist}(p', q)^2| \\ &\leq \text{dist}(p, q')^2 - \text{dist}(p', q')^2 \\ &= (\text{dist}(p, q') + \text{dist}(p', q')) (\text{dist}(p, q') - \text{dist}(p', q')) \\ &\leq 2\Delta \cdot \text{dist}(p, p'), \end{aligned} \quad (37)$$

where Δ is the diameter of the sample space (assuming that the centers also reside in the sample space). Hence, $\rho = 2\Delta$ for k -means.

• For PCA, $\text{cost}(p, x) = \text{dist}(p, xp)^2$, where $x = WW^T$ is the projection matrix, and W is the matrix consisting of the first l ($l < d$) principle components as columns. For any

data points $p, p' \in \mathbb{R}^d$, assuming WLOG that $\text{dist}(p, xp) \geq \text{dist}(p', xp')$, we have

$$\begin{aligned} & |\text{dist}(p, xp)^2 - \text{dist}(p', xp')^2| \\ &= (\text{dist}(p, xp) + \text{dist}(p', xp')) \cdot (\text{dist}(p, xp) - \text{dist}(p', xp')). \end{aligned} \quad (38)$$

The first factor in (38) is upper-bounded by 2Δ (Δ : diameter of sample space), as long as the projections xp, xp' reside in the sample space. The second factor is upper-bounded by

$$\begin{aligned} & \text{dist}(p, xp') - \text{dist}(p', xp') + \text{dist}(xp', xp) \\ & \leq \text{dist}(p, p') + \|x\|_2 \cdot \text{dist}(p, p'), \end{aligned} \quad (39)$$

which is due to triangle inequality and the sub-multiplicative property of ℓ_2 norm, i.e., $\text{dist}(xp', xp) = \|x(p' - p)\|_2 \leq \|x\|_2 \cdot \|p' - p\|_2$. As the principle components are mutually orthogonal unit vectors, we have $\|x\|_2 \leq l$. Plugging these into (38) gives

$$|\text{dist}(p, xp)^2 - \text{dist}(p', xp')^2| \leq 2\Delta(l+1)\text{dist}(p, p'), \quad (40)$$

i.e., $\rho = 2\Delta(l+1)$ for PCA.

• For SVM, the cost function is defined as: $\text{cost}(p, x) = \max(0, 1 - p_d(p_{1:d-1}^T x_{1:d-1} + x_d))$, where $p_{1:d-1} \in \mathbb{R}^{d-1}$ denotes the numerical portion of a data point p , and $p_d \in \{1, -1\}$ denotes its label. Consider two points p, p' that are identical with each other except the label, i.e., $p_{1:d-1} = p'_{1:d-1}$ and $p_d = -p'_d$. Suppose that $p_{1:d-1}^T x_{1:d-1} + x_d < -1$, $p_d = 1$, and $p'_d = -1$. Then we have $\text{dist}(p, p') = 2$, $\text{cost}(p', x) = 0$, and $\text{cost}(p, x) = 1 - (p_{1:d-1}^T x_{1:d-1} + x_d)$. As $\text{cost}(p, x)$ for SVM is unbounded in general, the ratio

$$\frac{|\text{cost}(p, x) - \text{cost}(p', x)|}{\text{dist}(p, p')} = \frac{1}{2} [1 - (p_{1:d-1}^T x_{1:d-1} + x_d)] \quad (41)$$

is unbounded. Therefore, $\rho = \infty$ for SVM.

APPENDIX C: ANALYSIS OF DIMENSION OF FUNCTION SPACE

Definition C.1 ([7]). Let $F := \{f_x : x \in \mathcal{X}\}$, where each element $f_x : P \rightarrow \mathbb{R}_{\geq 0}$ is a function from a set P to nonnegative real numbers. Define $B(x, r) := \{p \in P : f_x(p) \leq r\}$. The dimension of the function space $\dim(F, P)$ is the smallest integer m such that

$$|\{S \cap B(x, r) : x \in \mathcal{X}, r \geq 0\}| \leq |S|^m, \quad \forall S \subseteq P. \quad (42)$$

The dimension of function space is related to the Vapnik-Chervonenkis (VC) dimension, defined as follows. We refer to (P, \mathcal{R}) as a range space if P is a set of points and \mathcal{R} is a family of subsets of P .

Definition C.2 (VC dimension [51]). The VC dimension of a range space (P, \mathcal{R}) , denoted by $d_{\text{vc}}(P, \mathcal{R})$, is the maximum cardinality of a set $S \subseteq P$ that is shattered by \mathcal{R} , i.e., $\{S \cap R : R \in \mathcal{R}\}$ contains all the subsets of S .

Lemma C.1 (Corollary 5.2.3 [51]). If (P, \mathcal{R}) is a range space of VC dimension m , then for every $S \subseteq P$, we have $|\{S \cap R : R \in \mathcal{R}\}| \leq |S|^m$.

Table V
ERROR BOUND ϵ FOR MEB ($z = 1$)

dataset	coreset size	max relative error	ϵ
Fisher's iris	20	0.0053	0.1073
Facebook	80	0.0344	1.18
Pendigits	400	0.0026	2.1112
MNIST	400	0.0024	10.74
HAR	400	5.6054e-05	3.9212

By Definition C.1 and Lemma C.1, we have $\dim(F, P) \leq d_{\text{vc}}(P, \mathcal{R})$ for $\mathcal{R} := \{B(x, r) : x \in \mathcal{X}, r \geq 0\}$. This allows us to bound the dimension of a function space by bounding the VC dimension of the corresponding range space. The VC dimension has an intuitive meaning that it is the number of free parameters to uniquely specify a set in \mathcal{R} , e.g., the VC dimension of intervals is 2, the VC dimension of planar disks is 3, and the VC dimension of half spaces in \mathbb{R}^d is $d+1$ [51]. In our case, we conjecture that the VC dimension is $O(d_{\mathcal{X}})$, where $d_{\mathcal{X}}$ is the number of parameters to uniquely specify an $x \in \mathcal{X}$.

APPENDIX D: ADDITIONAL EVALUATIONS

D.1. Error Bound

In addition to the cost and accuracy, we have also evaluated the relative error in using the coreset to approximate the cost on the original dataset, defined as $|\text{cost}(P, x_S) - \text{cost}(S, x_S)| / \text{cost}(P, x_S)$, where x_S is the model parameter computed on a coreset S . By Definition II.1, this error should be upper-bounded by ϵ if S is an ϵ -coreset. In Tables V–VI, we show the maximum relative approximation error over all the Monte Carlo runs for MEB, together with the value of ϵ computed according to (9). Indeed, the error is always upper-bounded by ϵ . We note that ϵ tends to grow with the dimension of the dataset and can be greater than one. This is because the bound in (9) is based on the maximum Euclidean distance between a data point and the corresponding coreset point representing it, which tends to grow with the dimension of the dataset. Meanwhile, we also note that the bound tends to be loose, and the actual approximation error can be much smaller than ϵ . This is because by Definition II.1, ϵ needs to upper-bound the relative approximation error for any model parameter $x \in \mathcal{X}$, and may thus be loose at the computed model parameter x_S . We note that the bounds for the other machine learning problems (k -means, PCA, SVM, NN) equal the bound for MEB scaled by their corresponding Lipschitz constants, which can be large or even infinite (see Table II). Nevertheless, our experiments have shown that the proposed coresets can be used to train these models with competitive performances in both cost and accuracy.

D.2. Original Machine Learning Performance

We provide the performance of applying each machine learning algorithm directly to the original dataset in Table VII. For unsupervised learning problems (MEB, k -means, PCA), these were the denominators in evaluating the normalized costs of models learned on coresets. Each machine learning problem is parameterized as in Section V.

Table VI
ERROR BOUND ϵ FOR MEB ($z = 2$)

dataset	coreset size	max relative error	ϵ
Fisher's iris	20	1.1896e-05	0.1093
Facebook	80	1.9976e-06	1.3711
Pendigits	400	4.3876e-05	2.0257
MNIST	400	0.0020	8.53
HAR	400	4.9972e-07	3.5612

Table VII
ORIGINAL MACHINE LEARNING PERFORMANCE (COST FOR MEB, k -MEANS, PCA; ACCURACY FOR SVM AND NN)

problem	Fisher's	Facebook	Pendigits	MNIST	HAR
MEB	2.05	7.65	18.06	90.43	60.26
k -means	84.54	853.67	2.00e+05	4.98e+07	2.86e+06
PCA	1.94	197.19	910.79	1.73e+04	4.46e+05
SVM/NN ⁴	100%	89.36%	99.33%	87.01%	78.01%

D.3. Additional Evaluations in Distributed Setting

To evaluate the impact of the coreset size (N) and the maximum number of local centers (K), we vary each of these parameters while fixing the data distribution scheme as the ‘uniform’ scheme. Figures 15–16 are the results of varying N under fixed K . As expected, all the algorithms benefit from a larger coreset size, but we observe more significant improvements for the distributed algorithms (CDCC and DRCC). Figures 17–18 show the results of varying K under fixed N . Although RCC-kmeans and CDCC do not depend on K , we still show them as benchmarks. Overall, increasing K improves the quality of the models trained on the coreset generated by DRCC, as it gives DRCC more space to optimize its configuration. Specifically, we see that increasing K from 1 to 2 notably improves the performance in most cases, but further increasing K does not bring notable improvements for the unsupervised learning problems. The accuracy of NN, in contrast, keeps improving as K increases. We note that the small dip at $K = 2$ in the k -means cost is because this parameter leads to 2 local centers at most nodes, which coincides with the number of global centers we are computing.

REFERENCES

- [1] H. Lu, M. Li, T. He, S. Wang, V. Narayanan, and K. S. Chan, “Robust coreset construction for distributed machine learning,” in *2019 IEEE Global Communications Conference (GLOBECOM)*, 2019.
- [2] M. Chiang and T. Zhang, “Fog and IoT: An overview of research opportunities,” *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 854–864, December 2016.
- [3] D. Pietro-Barral and B. Guijarro-Berdinas, “A survey of methods for distributed machine learning,” in *Progress in Artificial Intelligence*, November 2012.
- [4] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *AISTATS*, May 2016.
- [5] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, “When edge meets learning: Adaptive control for resource-constrained distributed machine learning,” in *IEEE INFOCOM*, April 2018.

⁴As mentioned before, for Fisher’s iris, Facebook, and Pendigits, we evaluated SVM accuracy; for MNIST and HAR, we evaluated NN accuracy.

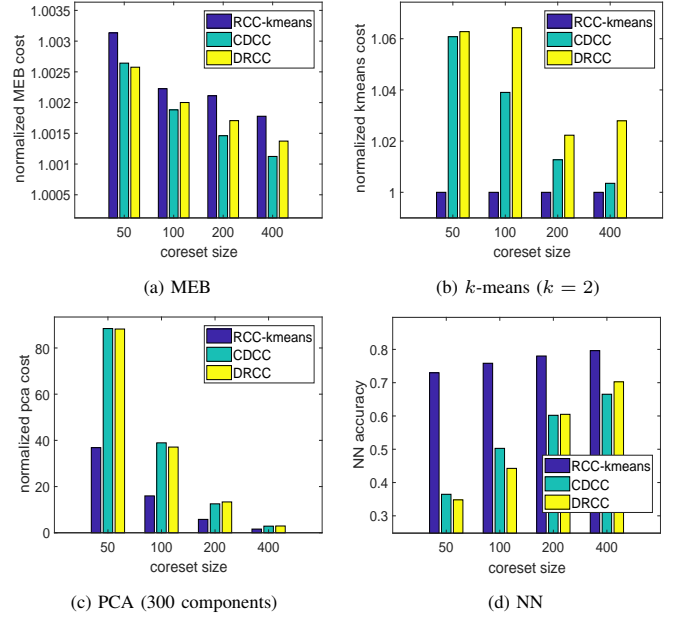


Figure 15. Evaluation on MNIST in distributed setting with varying coreset size ($K = 10$).

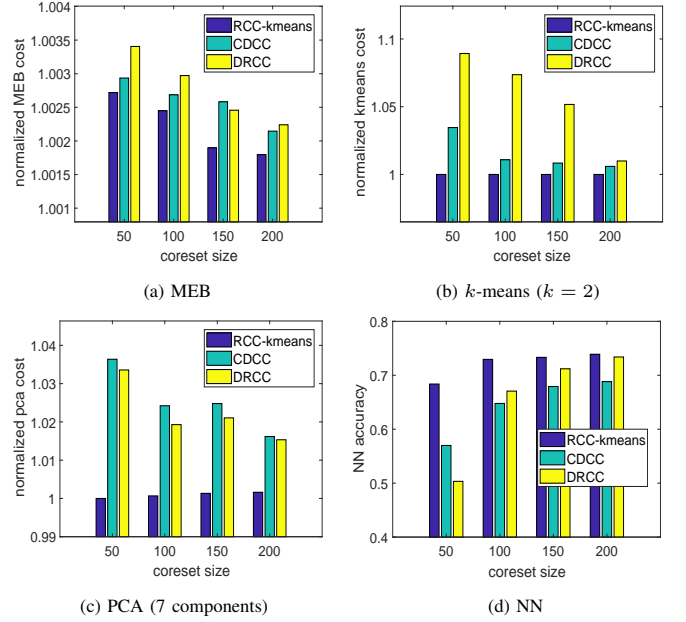


Figure 16. Evaluation on HAR in distributed setting with varying coreset size ($K = 10$).

- [6] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.
- [7] M. F. Balcan, S. Ehrlich, and Y. Liang, “Distributed k -means and k -median clustering on general topologies,” in *NIPS*, 2013.
- [8] R. Kannan, S. Vempala, and D. Woodruff, “Principal component analysis and higher correlations for distributed data,” in *COLT*, 2014.
- [9] A. Barger and D. Feldman, “ k -means for streaming and distributed big sparse data,” in *SDM*, 2016.
- [10] M. Bădoiu, S. Har-Peled, and P. Indyk, “Approximate clustering via core-sets,” in *ACM STOC*, 2002.
- [11] P. K. Chan and S. J. Stolfo, “Toward parallel and distributed learning by meta-learning,” in *AAAI Workshop in Knowledge Discovery in Databases*, 1997.
- [12] J. Kittler, M. Hatef, R. P. Duin, and J. Matas, “On combining classifiers,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*,

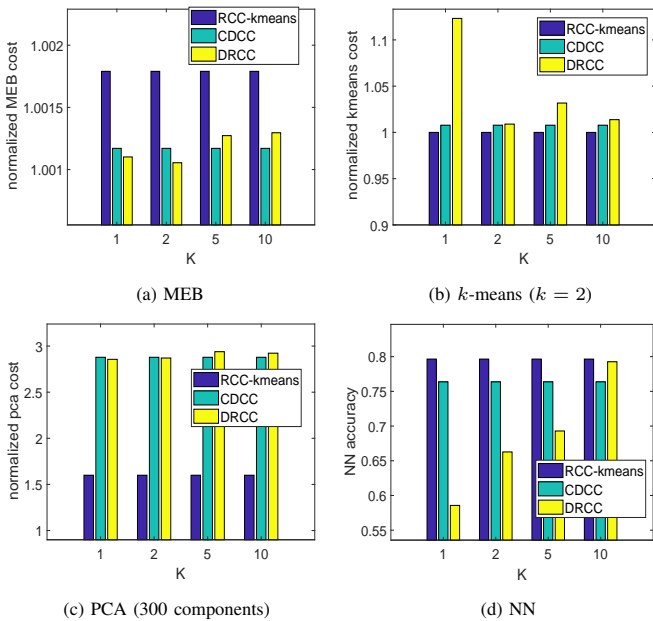


Figure 17. Evaluation on MNIST in distributed setting with varying K ($N = 400$).

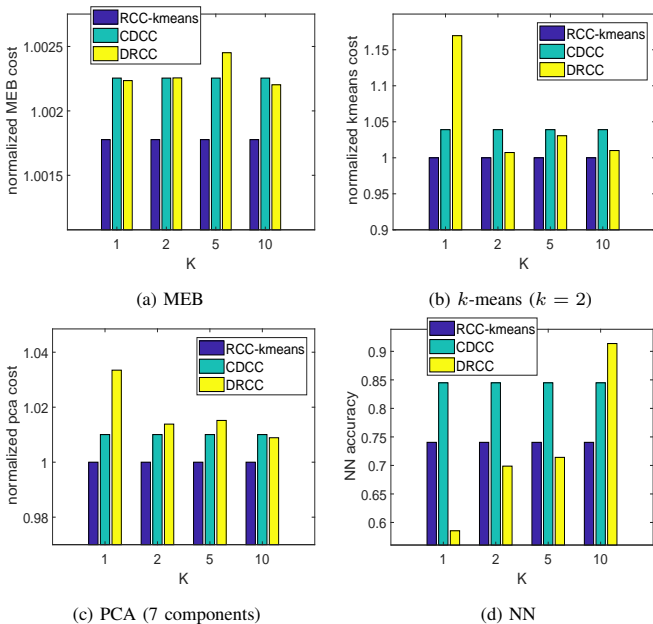


Figure 18. Evaluation on HAR in distributed setting with varying K ($N = 200$).


vol. 20, no. 3, pp. 226–239, March 1998.

- [13] D. Wolpert, “Stacked generalization,” *Neural Networks*, vol. 5, no. 2, pp. 241–259, 1992.
- [14] G. Tsoumakas and I. Vlahavas, “Effective stacking of distributed classifiers,” in *ECAI*, 2002.
- [15] Y. Guo and J. Sutiwaraphun, “Probing knowledge in distributed data mining,” in *Methodologies for Knowledge Discovery and Data Mining*, 1999.
- [16] N. Chawla, L. Halla, K. Bowyer, and W. Kegelmeyer, “Learning ensembles from bites: A scalable and accurate approach,” *Journal of Machine Learning Research*, vol. 5, pp. 421–451, April 2004.
- [17] A. Lazarevic and Z. Obradovic, “Boosting algorithms for parallel and distributed learning,” *Distributed and Parallel Databases*, vol. 11, no. 2, pp. 203–229, March 2002.
- [18] V. Smith, C.-K. Chiang, M. Sanjabi, and A. S. Talwalkar, “Federated multi-task learning,” in *Advances in Neural Information Processing Systems*, 2017, pp. 4424–4434.
- [19] J. M. Phillips, “Coresets and sketches,” *CoRR*, vol. abs/1601.00617, 2016.
- [20] A. Munteanu and C. Schwiiegelshohn, “Coresets-methods and history: A theoreticians design pattern for approximation and streaming algorithms,” *KI - Künstliche Intelligenz*, vol. 32, no. 1, pp. 37–53, 2018.
- [21] J. Wang, J. D. Lee, M. Mahdavi, M. Kolar, N. Srebro *et al.*, “Sketching meets random projection in the dual: A provable recovery algorithm for big and high-dimensional data,” *Electronic Journal of Statistics*, vol. 11, no. 2, pp. 4896–4944, 2017.
- [22] K. Makarychev, Y. Makarychev, and I. Razenshteyn, “Performance of johnson-lindenstrauss transform for k-means and k-medians clustering,” in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 2019, pp. 1027–1038.
- [23] D. Feldman and M. Langberg, “A unified framework for approximating and clustering data,” in *STOC*, June 2011.
- [24] M. Badoiu and K. L. Clarkson, “Smaller core-sets for balls,” in *SODA*, 2003.
- [25] S. Har-Peled and K. R. Varadarajan, “Projective clustering in high dimensions using core-sets,” in *SOCC*, 2002.
- [26] I. W. Tsang, J. T. Kwok, and P.-M. Cheung, “Core vector machines: Fast SVM training on very large data sets,” *The Journal of Machine Learning Research*, vol. 6, pp. 363–392, December 2005.
- [27] S. Har-Peled, D. Roth, and D. Zimak, “Maximum margin coresets for active and noise tolerant learning,” in *IJCAI*, 2007.
- [28] D. Feldman, M. Volkov, and D. Rus, “Dimensionality reduction of massive sparse datasets using coresets,” in *NIPS*, 2016.
- [29] D. Feldman, A. Munteanu, and C. Sohler, “Smallest enclosing ball for probabilistic data,” in *SOCC*, 2014.
- [30] K. L. Clarkson, “Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm,” *ACM Transactions on Algorithms*, vol. 6, no. 4, August 2010.
- [31] M. Langberg and L. J. Schulman, “Universal ϵ approximators for integrals,” in *SODA*, 2010.
- [32] D. Feldman, M. Feigin, and N. Sochen, “Learning big (image) data via coresets for dictionaries,” *Journal of Mathematical Imaging and Vision*, vol. 46, no. 3, pp. 276–291, March 2013.
- [33] A. Molina, A. Munteanu, and K. Kersting, “Core dependency networks,” in *AAAI*, 2018.
- [34] V. Braverman, D. Feldman, and H. Lang, “New frameworks for offline and streaming coreset constructions,” *CoRR*, vol. abs/1612.00889, 2016.
- [35] D. Feldman, A. Fiat, and M. Sharir, “Coresets for weighted facilities and their applications,” in *FOCS*, 2006.
- [36] G. Frahling and C. Sohler, “Coresets in dynamic geometric data streams,” in *STOC*, 2005.
- [37] S. Har-Peled and S. Mazumdar, “On coresets for k-means and k-median clustering,” in *STOC*, 2004.
- [38] D. Feldman, M. Schmidt, and C. Sohler, “Turning big data into tiny data: Constant-size coresets for k-means, PCA, and projective clustering,” in *SODA*, 2013.
- [39] C. Boutsidis, P. Drineas, and M. Magdon-Ismail, “Near-optimal coresets for least-squares regression,” *IEEE Trans. IT*, vol. 59, no. 10, pp. 6880–6892, October 2013.
- [40] D. Feldman, A. Krause, and M. Faulkner, “Scalable training of mixture models via coresets,” in *NIPS*, 2011.
- [41] D. Aloise, A. Deshpande, P. Hansen, and P. Popat, “NP-hardness of Euclidean sum-of-squares clustering,” *Machine Learning*, vol. 75, no. 2, pp. 245–248, May 2009.
- [42] N. Megiddo and K. J. Supowit, “On the complexity of some common geometric location problems,” *SIAM Journal of Computing*, vol. 13, no. 1, pp. 182–196, 1984.
- [43] D. Arthur and S. Vassilvitskii, “k-means++: The advantages of careful seeding,” in *SODA*, January 2007.
- [44] M. Cohen, Y. T. Lee, G. Miller, J. Pachocki, and A. Sidford, “Geometric median in nearly linear time,” in *STOC*, 2016.
- [45] A. Virmaux and K. Scaman, “Lipschitz regularity of deep neural networks: analysis and efficient estimation,” in *Advances in Neural Information Processing Systems*, 2018, pp. 3835–3844.
- [46] R. Fisher, “Iris data set,” <https://archive.ics.uci.edu/ml/datasets/iris>, 1936. [Online]. Available: <https://archive.ics.uci.edu/ml/datasets/iris>
- [47] S. Moro, P. Rita, and B. Vala, “Facebook metrics data set,” <https://archive.ics.uci.edu/ml/datasets/Facebook+metrics>, 2016. [Online]. Available: <https://archive.ics.uci.edu/ml/datasets/Facebook+metrics>
- [48] E. Alpaydin and F. Alimoglu, “Pen-based recognition of handwritten digits data set,” <https://archive.ics.uci.edu/ml/datasets/Pen-Based+Recognition+of+Handwritten+Digits>, 1996. [Online]. Available: <https://archive.ics.uci.edu/ml/datasets/Pen-Based+Recognition+of+Handwritten+Digits>


- [49] Y. LeCun, C. Cortes, and C. Burges, "The MNIST database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, 1998. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [50] D. Anguita, A. Ghio, L. Oneto, X. Parra, and J. L. Reyes-Ortiz, "A public domain dataset for human activity recognition using smartphones," in *ESANN*, 2013.
- [51] S. Har-Peled, *Geometric Approximation Algorithms*. American Mathematical Society, 2011.




Hanlin Lu (S'19) received the B.S. degree in mathematics from Beihang University in 2017. He is a Ph.D. candidate in Computer Science and Engineering at Pennsylvania State University, advised by Prof. Ting He. His research interest includes coresets and distributed machine learning.



Ming-Ju Li is a Master's student in Electrical and Computer Engineering as well as a research assistant in Network Science Research Group under the instruction of Dr. Ting He at The Pennsylvania State University. His research interests are Computer Network and Machine Learning.



Ting He (SM'13) received the B.S. degree in computer science from Peking University, China, in 2003 and the Ph.D. degree in electrical and computer engineering from Cornell University, Ithaca, NY, in 2007. Dr. He is an Associate Professor in the School of Electrical Engineering and Computer Science at Pennsylvania State University, University Park, PA. Between 2007 and 2016, she was a Research Staff Member in the Network Analytics Research Group at the IBM T.J. Watson Research Center, Yorktown Heights, NY. Her work is in the broad areas of computer networking, network modeling and optimization, and statistical inference. Dr. He is an Associate Editor for *IEEE Transactions on Communications* (2017-2020) and *IEEE/ACM Transactions on Networking* (2017-2021). She was the Membership co-chair of ACM N2Women in 2013-2014 and was listed in "N2Women: Rising Stars in Networking and Communications" in 2017. She received the Research Division Award and multiple Outstanding Contributor Awards from IBM, the Most Collaboratively Complete Publications Award from ITA, the Best Paper Award at the 2013 International Conference on Distributed Computing Systems (ICDCS), the Outstanding Student Paper Award at the 2015 ACM SIGMETRICS, and the Best Student Paper Award at the 2005 International Conference on Acoustic, Speech and Signal Processing (ICASSP).



Shiqiang Wang (S'13–M'15) received his Ph.D. from the Department of Electrical and Electronic Engineering, Imperial College London, United Kingdom, in 2015. Before that, he received his master's and bachelor's degrees at Northeastern University, China, in 2011 and 2009, respectively. He joined IBM T. J. Watson Research Center in 2016 as a Research Staff Member, where he was also a Graduate-level Co-op in the summers of 2014 and 2013. In the fall of 2012, he was at NEC Laboratories Europe, Heidelberg, Germany. His current research focuses on theoretical and practical aspects of mobile edge computing, cloud computing, and machine learning. Dr. Wang serves as an associate editor of *IEEE Access*. He served as a technical program committee (TPC) member of several international conferences including IEEE ICDCS, IJCAI, WWW, IFIP Networking, IEEE GLOBECOM, IEEE ICC, and as a reviewer for a number of international journals and conferences. He received the IBM Outstanding Technical Achievement Award (OTAA) in 2019, multiple Invention Achievement Awards from IBM since 2016, Best Paper Finalist of the IEEE International Conference on Image Processing (ICIP) 2019, and Best Student Paper Award of the Network and Information Sciences International Technology Alliance (NIS-ITA) in 2015.



Vijaykrishnan Narayanan (F'11) is the Robert Noll Chair of Computer Science and Engineering and Electrical Engineering at The Pennsylvania State University. His research interests are in Computer Architecture, Power-Aware Computing, and Intelligent Embedded Systems. He is a Fellow of IEEE and ACM.



Kevin Chan (S'02–M'09–SM'18) received the B.S. degree in electrical and computer engineering and engineering and public policy from Carnegie Mellon University, Pittsburgh, PA, USA, in 2001, and the M.S. and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2003 and 2008, respectively. He is currently a Research Scientist with the Computational and Information Sciences Directorate, U.S. Army Combat Capabilities Development Command, Army Research Laboratory, Adelphi, MD, USA. He is actively involved in research on network science, distributed analytics, and cybersecurity. He has received multiple best paper awards and the NATO Scientific Achievement Award. He has served on the Technical Program Committee for several international conferences, including IEEE DCSS, IEEE SECON, and IEEE MILCOM. He is a Co-Editor of the *IEEE Communications Magazine Military Communications and Networks Series*.