
Rk-means: Fast Clustering for Relational Data

Ryan R. Curtin
RelationalAI
ryan@ratml.org

Benjamin Moseley
Tepper School of Business
Carnegie Mellon University
moseleyb@andrew.cmu.edu

Hung Q. Ngo
RelationalAI
hung.q.ngo@relational.ai

XuanLong Nguyen
Department of Statistics
University of Michigan
xuanlong@umich.edu

Dan Olteanu
University of Oxford
dan.olteanu@cs.ox.ac.uk

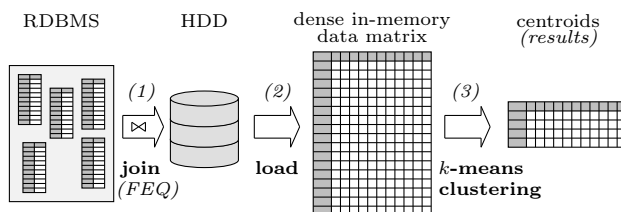
Maximilian Schleich
University of Oxford
max.schleich@cs.ox.ac.uk

Abstract


Conventional machine learning algorithms cannot be applied until a data matrix is available to process. When the data matrix needs to be obtained from a relational database via a feature extraction query, the computation cost can be prohibitive, as the data matrix may be (much) larger than the total input relation size. This paper introduces Rk-means, or relational k -means algorithm, for clustering relational data tuples without having to access the full data matrix. As such, we avoid having to run the expensive feature extraction query and storing its output. Our algorithm leverages the underlying structures in relational data. It involves construction of a small *grid coreset* of the data matrix for subsequent cluster construction. This gives a constant approximation for the k -means objective, while having asymptotic runtime improvements over standard approaches of first running the database query and then clustering. Empirical results show orders-of-magnitude speedup, and Rk-means can run faster on the database than even just computing the data matrix.

1 Introduction

Clustering is an ubiquitous technique for exploratory data analysis, whether applied to small samples or industrial-scale data. In the latter setting, two steps are typically performed: (1) *data preparation*, or extract-



(a) Typical k -means data science workflow. Alternate representations can be used for the data in step (2) for greater computational efficiency (e.g., streaming); and, approximation strategies are known for step (3). However, the dataset often comes from an underlying database system, and in this case the expensive FEQ join (1) is unavoidable.



(b) The Rk-means data science workflow. We avoid ever computing the expensive FEQ by instead clustering each underlying relation (steps 1 and 2, Section 4); we then use FAQs for efficient weighted k -means of the cross-product of those relations (steps 3 and 4, Section 4). This gives significant empirical and theoretical accelerations, and bounded approximation—*without ever computing the full data matrix*.

Figure 1: Conventional k -means and Rk-means.

transform-load (ETL) operations, and (2) *clustering the extracted data*—often with a technique such as the popular k -means algorithm (Cady, 2017; Wu et al., 2008). In this setting, data often reside in a relational database, requiring a *feature extraction query* (FEQ) to be performed, *joining* involved relations together to form the data matrix: each row corresponds to a data tuple and each column a feature. Then, the data matrix is used as input to a clustering algorithm. This matrix can be expensive to compute, and may take up

space asymptotically larger than the database itself. Moreover, the join computation time may exceed the time it takes to obtain clusters. As a result it is not uncommon that the exploratory trip into the dataset may be stopped right at the gate.

As an example, consider a retailer database consisting of three tables: `product`, which contains data about p products, `store`, which contains data about s stores, and `transaction`, which contains the number of transactions for each (product, store) combination on a given day. A practitioner may want to cluster each (product, store) combination as part of an analysis to determine items with related sales patterns across different stores for a given week. To do this, she constructs a data matrix containing all (product, store) combinations (including those with zero sales) for a given week, and additional attributes for each product and store. This can be achieved, for instance, by the following feature extraction query, given in SQLite syntax:

```
SELECT P.id AS i, S.id AS s, P.type AS t, P.price AS p,
       S.yelp_rating AS y, sum(ifnull(T.count, 0)) AS c
FROM product P, store S LEFT JOIN transactions T
ON T.product_id == P.id AND T.store_id == S.id
AND T.date BETWEEN '2019-05-13' AND '2019-05-20'
GROUP BY P.id, S.id;
```

The result of this query is of size $\Theta(ps)$. But the `transaction` table can be significantly smaller than this, since many stores may have zero sales of a particular product in a given week. Thus, the size of the data matrix can be asymptotically greater than the total input relations’ sizes. Real-world FEQs possess a similar explosion in both space and time complexity, only at a much larger scale, since they generally involve many more aggregations and tables. In Section 5, we present a real dataset from a large US retailer. The database has 6 tables of total size 1.5GB. The FEQ result, however, takes up 18GB, and constructing it takes longer than running a learning algorithm on it.

Stripping away the language of databases, a fundamental challenge is how to learn about the joint distribution of a data population given only marginal samples revealed by relational tables. This is possible when the objective function of an underlying model admits some factorization structure similar to conditional independence in graphical models (Koller and Friedman, 2009). This insight was exploited recently by database theorists to devise algorithms evaluating a generic class of relational queries called *functional aggregate queries*, or FAQs (Abo Khamis et al., 2016). The ability to answer FAQs quickly is a building block for a new class of efficient algorithms for training supervised learning models over relational data, *without* having to materialize (i.e. compute) the entire data matrix (Abo Khamis et al., 2018b,a; Schleich et al., 2016).

The goal of this paper is to devise a method for fast clus-

tering of relational data, without having to materialize the full data matrix. The challenge of unsupervised learning tasks in general and the k -means algorithm in particular is that the learning objective is not decomposable across marginal samples in relational tables. To enable fast relational computation, we utilize the idea of constructing a *grid coreset*—a small set of points that provide a good summarization of the original (and unmaterialized) data tuples, based on which a provably good clustering can be obtained.

The resulting algorithm, which we call Rk-means, has several remarkable properties. First, Rk-means has a provable constant approximation guarantee relative to the k -means objective, despite the fact that the algorithm does not require access to the full data matrix. Our approximation analysis is established via a connection of Rk-means to the theory of optimal transport (Villani, 2009). Second, Rk-means is enhanced by leveraging structures prevalent in relational data: categorical variables, functional dependencies, and the topology of feature extraction queries. These structures lead to exponential reduction in coreset size without incurring loss in the coreset’s approximation error. We show that Rk-means is provably more efficient both in time and space complexity when comparing against the application of the vanilla k -means to the full data matrix. Finally, experimental results show significant speedups with little loss in the k -means objective. We observe orders-of-magnitude improvement in the running time compared to traditional methods. Rk-means operates fine when other approaches would run out of memory, enabling clustering on truly massive datasets.

2 Discussion of related work

Coresets for clustering. From early work on k -means algorithm (Lloyd, 1982), ideas emerged for acceleration via coresets (Har-Peled and Mazumdar, 2004; Bachem et al., 2018). Coresets have become the cornerstone of modern streaming algorithms (Guha et al., 2003; Braverman et al., 2017), massively parallel (MPC) algorithms (Ene et al., 2011; Bahmani et al., 2012), and are used to speed up sequential methods (Meyerson et al., 2004; Sohler and Woodruff, 2018).

Unfortunately, existing algorithms for coreset construction do not readily lend themselves to the relational setting; there are several hurdles. First of all, coresets are formed by constructing a set S of data points (tuples) that represent the entire data set X well. Typically, S is a *weighted* representation of the data, where each point in the universe contributes one unit of weight to its closest point in X (Har-Peled and Mazumdar, 2018; Ene et al., 2011; Balcan et al., 2013). In our relational setting, X can *only* be formed by computing the FEQ, but our goal is to *avoid* materializing X .

Unfortunately, for our setting, most existing coresets algorithms construct S in phases by determining the farthest points from S (Thorup, 2001; Arthur and Vassilvitskii, 2007; Har-Peled and Mazumdar, 2018). This is difficult without X fully materialized. Another difficulty is that, even if the points in S are given, weighting the points in S is an open problem for relational algorithms (Khamis et al., 2019): consider that for a point $x \in S$, finding the number of closest points in (unmaterialized) X is non-trivial, as the points and their attributes are stored across several tables. No method, either deterministic or stochastic (e.g., sampling), is known that runs in time asymptotically faster than computing/materializing X . Our method avoids this by constructing a *grid coreset* S which can be decomposed over the tables in such a way that computing the weights of the points is a straightforward task.

Other Work. Our work draws inspiration from three lines of existing work and ideas: coresets for clustering (above), database algorithms, and optimal transport. As one example, database and disk hardware optimizations have been considered to improve clustering relational data (Ordonez, 2006; Ordonez and Omiecinski, 2004). Recent advances include the work of Abo Khamis et al. (2018b) and Schleich et al. (2016). k -means has also been connected to optimal transport, which goes back to at least (Pollard, 1982) (see also (Graf and Luschgy, 2000)). Recently this connection has received increased interest in the statistics and machine learning communities, resulting in fresh new clustering techniques (del Barrio et al., 2017; Ho et al., 2017; Ye et al., 2017). To our knowledge, these related lines of work have not been explored together. Motivated by clustering relational data, our attempt at solving a clustering problem formulated as optimal transport in the marginal (projected) spaces to scalably perform k -means clustering appears to be the first.

Finally, it is worth noting that despite its popularity, the basic k -means technique is not always a preferred choice in clustering categorical or high-dimensional data. One may either adopt other clustering techniques (Hartigan, 1975; Kaufman and Rousseeuw, 1990; Everitt et al., 2011), or modify the basic k -means method, e.g., by suitably placing weights on different features of mixed data types and replacing metric ℓ_2 by ℓ_p (Huang, 1998), or incorporating a regularizer to combat high dimensionality (Witten and Tibshirani, 2010; Sun et al., 2012). As we shall see, the relational techniques and associated theory that we introduce for the basic k -means extend easily to such improvements.

3 The Rk-means algorithm

Although the Rk-means algorithm is motivated by application to relational databases, its basic idea is also

Algorithm 1 Rk-means: k -means via grid-coreset

- 1: **Input:** query Q , number of clusters k
 - 2: **Input:** $[d] = S_1 \cup \dots \cup S_m$, $\kappa \geq 2$
 - 3: **Output:** centroids $\mathbf{C} \in \mathcal{R}^{k \times d}$
 - 4: **for** $j = 1$ to m **do**
 - 5: $\mathbf{X}_j \leftarrow \{\mathbf{x}_{S_j} \mid \mathbf{x} \in \mathbf{X}\}$
 - 6: $w_j \leftarrow$ weight function defined in (2)
 - 7: $\mathbf{C}_j \leftarrow \text{wkmeans}_1(\mathbf{X}_j, w_j, \kappa)$
 - 8: **end for**
 - 9: $\mathbf{G} \leftarrow \mathbf{C}_1 \times \dots \times \mathbf{C}_m$ {the grid coreset}
 - 10: $w_{\text{grid}} \leftarrow$ weight function defined in (3)
 - 11: $\mathbf{C} \leftarrow \text{wkmeans}_2(\mathbf{G}, w_{\text{grid}}, k)$
-

of independent interest and can be easily described without the database language.

First we define the *weighted k -means* problem, which Rk-means solves (weights are also handy in combining mixed data types (Huang, 1998)). Let \mathbf{X} be a set of points in \mathbb{R}^d , and \mathbf{Y} be a non-empty set of points in the same space. Let $d(\mathbf{x}, \mathbf{Y}) := \min_{\mathbf{y} \in \mathbf{Y}} \|\mathbf{x} - \mathbf{y}\|$ denote the minimum distance from \mathbf{x} to an element in \mathbf{Y} . In some cases, the ℓ_2 norm $\|\cdot\|$ may be replaced by the ℓ_p norm $\|\cdot\|_p$ for some $p \geq 1$. A *weighted k -means instance* is a pair (\mathbf{X}, w) , where \mathbf{X} is a set of points in \mathbb{R}^d and $w : \mathbf{X} \rightarrow \mathbb{R}^+$ is a weight function. Without loss of generality, assume $\sum_{\mathbf{x} \in \mathbf{X}} w(\mathbf{x}) = 1$. The task is to find a set $\mathbf{C} = \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_k\}$ of k centroids to minimize the objective $L(\mathbf{X}, \mathbf{C}, w) = \sum_{\mathbf{x} \in \mathbf{X}} w(\mathbf{x})d(\mathbf{x}, \mathbf{C})^2$.

That is, we want to solve the problem $\text{OPT}(\mathbf{X}, w) := \min_{\mathbf{C}} L(\mathbf{X}, \mathbf{C}, w)$. With Rk-means, we will do this by projecting \mathbf{X} onto different sets of coordinates, and clustering each projection individually. To this end, let $[d] = S_1 \cup \dots \cup S_m$ denote an arbitrary *partition* of the dimensions $[d]$ into non-empty subsets. For every $\mathbf{x} \in \mathbb{R}^d$, and $j \in [m]$, let \mathbf{x}_{S_j} denote the projection of \mathbf{x} onto the coordinates in S_j . Define the projection set \mathbf{X}_j and corresponding weight function $w_j : \mathbb{R}^{S_j} \rightarrow \mathbb{R}$ by

$$\mathbf{X}_j := \{\mathbf{x}_{S_j} \mid \mathbf{x} \in \mathbf{X}\}, \quad (1)$$

$$w_j(\mathbf{z}) := \sum_{\mathbf{x} \in \mathbf{X} : \mathbf{x}_{S_j} = \mathbf{z}} w(\mathbf{x}). \quad (2)$$

In words, the w_j are the *marginal measures* of w on the subspace of coordinates S_j . With these notations established, Algorithm 1 presents the high-level description of our algorithm, Rk-means.

For each $j \in [m]$, in line 7 we perform k -means to obtain κ individual clusters on each subspace S_j for some $\kappa \geq 2$. These are solved using some weighted k -means algorithm denoted by wkmeans_1 with approximation ratio α .¹ Then, using the results of these clusterings, we

¹That is, wkmeans_1 finds clusterings $\hat{\mathbf{C}}$ where $L(\mathbf{X}, \hat{\mathbf{C}}, w)/L(\mathbf{X}, \mathbf{C}^*, w) \leq \alpha$, where $\mathbf{C}^* = \text{argmin}_{\mathbf{C}} L(\mathbf{X}, \mathbf{C}, w)$.

assemble a cross-product weighted grid \mathbf{G} of centroids, and then perform k -means clustering on these using the algorithm denoted wkmeans_2 to reduce down to the desired result of k centroids. Typically, we take $\kappa = O(k)$. Note also that Rk-means could easily be modified to use a different κ value for different subspaces S_j .

Let $\mathbf{X} := \bigsqcup_{\mathbf{g} \in \mathbf{G}} \mathbf{X}_{\mathbf{g}}$ denote a partition of \mathbf{X} into $|\mathbf{G}|$ parts, where $\mathbf{X}_{\mathbf{g}}$ denote the set of points in \mathbf{X} closer to \mathbf{g} than other grid points in \mathbf{G} (breaking ties arbitrarily). Then, the weight function $w_{\text{grid}} : \mathbf{G} \rightarrow \mathbb{R}^+$ (line 11) is

$$w_{\text{grid}}(\mathbf{g}) := \sum_{\mathbf{x} \in \mathbf{X}_{\mathbf{g}}} w(\mathbf{x}). \quad (3)$$

Weighted k -means and optimal transport. We will analyze Rk-means in the language of optimal transport. The connection of k -means in general, and of our algorithm to optimal transport in particular, provides another interesting insight into our algorithm.

The *optimal transport distance* characterizes the distance between two probability measures, by measuring the optimal cost of transporting mass from one to another (Villani, 2009). Although this is defined more generally for any two probability measures in abstract spaces, for our purpose it is convenient to consider two discrete probability measures P and P' on \mathbb{R}^d .

Let \mathbf{Z} and \mathbf{Z}' be two finite point sets in \mathbb{R}^d . Let δ denote the Dirac measure. Let $P := \sum_{\mathbf{z} \in \mathbf{Z}} p(\mathbf{z})\delta_{\mathbf{z}}$ and $P' := \sum_{\mathbf{z}' \in \mathbf{Z}'} p'(\mathbf{z}')\delta_{\mathbf{z}'}$ be two measures with supports \mathbf{Z} and \mathbf{Z}' , respectively. The mass transportation plan is formalized by a *coupling*: a joint distribution $\mathbf{Q} = (q(\mathbf{z}, \mathbf{z}'))_{(\mathbf{z}, \mathbf{z}') \in \mathbf{Z} \times \mathbf{Z}'}$, where the marginal constraints $\sum_{\mathbf{z} \in \mathbf{Z}} q_{\mathbf{z}, \mathbf{z}'} = p'(\mathbf{z}')$ and $\sum_{\mathbf{z}' \in \mathbf{Z}'} q_{\mathbf{z}, \mathbf{z}'} = p(\mathbf{z})$ hold.

Definition 3.1. For any $p \geq 1$, the *Wasserstein distance* of order p is defined by the minimization of \mathbf{Q} over all possible couplings: $W_p(P, P') = \min_{\mathbf{Q}} \{\sum_{\mathbf{z}, \mathbf{z}'} q(\mathbf{z}, \mathbf{z}') \|\mathbf{z} - \mathbf{z}'\|_p^p\}^{1/p}$.

Let $P^{\text{in}} = \sum_{\mathbf{x} \in \mathbf{X}} w(\mathbf{x})\delta_{\mathbf{x}}$ be the discrete measure associated with the input instance of our weighted k -means problem; then, this can be expressed precisely as an optimal transport problem: $M^* = \arg \min W_2^2(M, P^{\text{in}})$, where the optimization is over the space of discrete measures M that have k support points (the set \mathbf{C} of k centroids). Note that $\text{OPT}(\mathbf{X}, w) = W_2^2(M^*, P^{\text{in}})$. Replacing ℓ_2 by say ℓ_1 , we obtain the k -median problem, for which the objective becomes $W_1(M^*, P^{\text{in}})$.

Approximation Analysis. We next analyze the approximation ratio of Rk-means working with the W_2^2 objective, provided that wkmeans_1 has approximation ratio α and wkmeans_2 has approximation ratio γ .² The reason we might want to invoke different algorithms to

solve these sub-problems is because, as we shall show in the next section, we may want to exploit the (relational) structures of the FEQ to construct a “nice” partition $S_1 \cup \dots \cup S_m$. We show that the overall approximation ratio of Rk-means is $(\sqrt{\alpha} + \sqrt{\gamma} + \sqrt{\alpha\gamma})^2$. In many common cases, the database has structure that allows $\alpha = 1$, yielding an overall ratio of $(1 + 2\sqrt{\gamma})^2$.

For our analysis it is useful to understand Algorithm 1 in the language of optimal transport. For any finite point set $\mathbf{Y} \subset \mathbb{R}^d$ and a measure $M = \sum_{\mathbf{y} \in \mathbf{Y}} p(\mathbf{y})\delta_{\mathbf{y}}$ with support \mathbf{Y} , define the marginal measures M_j on coordinates S_j induced by M in the natural way, i.e. $M_j := \sum_{\mathbf{z} \in \mathbf{Z}} p_j(\mathbf{z})\delta_{\mathbf{z}}$ where p_j is defined analogous to w_j in (2). Under this notation, P^{in} induces the marginal measures $P_j^{\text{in}} := \sum_{\mathbf{z} \in \mathbf{X}_j} w_j(\mathbf{z})\delta_{\mathbf{z}}$. Then, Algorithm 1 can be described by the following steps:

- (1) For each $j \in [m]$, pick M_j to be the (α -approximate) minimizer of $W_2^2(M_j, P_j^{\text{in}})$, where $\mathbf{C}_j = \text{supp}(M_j)$ is the support of M_j and $|\mathbf{C}_j| = \kappa$ (line 7).
- (2) Collect the κ^d grid points \mathbf{G} and let probability measure Q be the one with support in \mathbf{G} such that Q minimizes $W_2^2(Q, P^{\text{in}})$. (We solve this problem exactly!)
- (3) Finally, return P which is the measure with exactly k support points in \mathbb{R}^d that (γ -approximately) minimizes $W_2^2(P, Q)$ (line 11).

This is precisely the solution obtained by Algorithm 1. We present next some useful facts.

Lemma 3.2. For any discrete measure M on \mathbb{R}^d , $W_2^2(M, P^{\text{in}}) \geq \sum_{j=1}^m W_2^2(M_j, P_j^{\text{in}})$.

Proof. Immediate: a coupling of two measures induces valid marginal couplings of marginal measures. \square

Proposition 3.3. (a) If $\kappa \geq |\text{supp}(M_j^*)| \forall j \in [m]$, then $W_2(P^{\text{in}}, P) \leq (\sqrt{\gamma} + \sqrt{\alpha} + \sqrt{\alpha\gamma})W_2(P^{\text{in}}, M^*)$. (b) For any $\kappa \geq 1$, there exists a distribution P^{in} such that $\frac{W_2(P^{\text{in}}, P)}{W_2(P^{\text{in}}, M^*)} \geq \sqrt{1 - e^{-m/(2\kappa)}} \frac{\sqrt{3}k^{3/(2m)}}{2\kappa m^{1/2}}$.

The condition of part (a) is satisfied, for instance, by setting $\kappa = k$. In practice, $\kappa < k$ may suffice. Moreover, part (b) dictates that κ must grow with k appropriately for our algorithm to maintain a constant approximation guarantee. Since solution \mathbf{C} has cost $L(\mathbf{X}, \mathbf{C}, w) = W_2^2(P, P^{\text{in}})$, and $\text{OPT}(\mathbf{X}, w) = W_2^2(M^*, P^{\text{in}})$, the following theorem is immediate from Prop. 3.3(a).

Theorem 3.4. Suppose wkmeans_1 and wkmeans_2 have approximation ratios α and γ . Then by choosing $\kappa = k$, the solution \mathbf{C} given by Rk-means has the following guarantee: $L(\mathbf{X}, \mathbf{C}, w) \leq (\sqrt{\gamma} + \sqrt{\alpha} + \sqrt{\alpha\gamma})^2 \text{OPT}(\mathbf{X}, w)$.

Specifically, if both sub-problems are solved optimally ($\alpha = \gamma = 1$), Rk-means is a 9-approximation.

²The best known approximation ratio is 6.357 for data in Euclidean space (Ahmadian et al., 2017).

Proof of Prop. 3.3. (a) By the definition of Q , the optimal transport plan from P^{in} to Q is such that each support point $s \in S$ is received by all $x \in \mathbf{X}$ nearest to s compared to other points in S . So,

$$\begin{aligned} & W_2^2(P^{\text{in}}, Q) \\ &= \sum_{\mathbf{x} \in \mathbf{X}} w(\mathbf{x}) d(\mathbf{x}, \mathbf{G})^2 = \sum_{\mathbf{x} \in \mathbf{X}} w(\mathbf{x}) \sum_{j=1}^m d(\pi_{S_j} \mathbf{x}, \mathbf{C}_j)^2 \\ &= \sum_{j=1}^m \sum_{\mathbf{z} \in \mathbf{X}_j} w_j(\mathbf{z}) d(\mathbf{z}, \mathbf{C}_j)^2 = \sum_{j=1}^m W_2^2(M_j, P_j^{\text{in}}) \\ &\leq \alpha \sum_{j=1}^m W_2^2(M_j^*, P_j^{\text{in}}) \leq \alpha \cdot W_2^2(M^*, P^{\text{in}}). \end{aligned}$$

The second to last inequality is due to the α -approximation of wkmeans_1 , and condition that $|\text{supp}(M_j)| \geq |\text{supp}(M_j^*)|$. The last inequality follows from Lemma 3.2. Apply triangle inequality of W_2 (cf. Villani (2009), pg. 106).

$$\begin{aligned} & W_2(P^{\text{in}}, P) & (4) \\ & \leq W_2(P^{\text{in}}, Q) + W_2(Q, P) & (5) \\ & \leq W_2(P^{\text{in}}, Q) + \sqrt{\gamma} \cdot W_2(Q, M^*) & (6) \\ & \leq W_2(P^{\text{in}}, Q) + \sqrt{\gamma} (W_2(Q, P^{\text{in}}) + W_2(P^{\text{in}}, M^*)) & (7) \\ & = (1 + \sqrt{\gamma}) W_2(P^{\text{in}}, Q) + \sqrt{\gamma} \cdot W_2(P^{\text{in}}, M^*) & (8) \\ & \leq (1 + \sqrt{\gamma}) \sqrt{\alpha} W_2(P^{\text{in}}, M^*) + \sqrt{\gamma} \cdot W_2(P^{\text{in}}, M^*) & (9) \\ & = (\sqrt{\alpha} + \sqrt{\gamma} + \sqrt{\alpha\gamma}) \cdot W_2(M^*, P^{\text{in}}). & (10) \end{aligned}$$

The second inequality is due to the fact that wkmeans_2 has approximation ratio γ ; the first and third are again triangle inequalities. We conclude the proof.

(b) We need only construct an example of P^{in} for the case $d = m$. Although P^{in} as an input to the algorithm is a discrete measure, for our purposes it suffices to take P^{in} to be the uniform distribution on $[0, 1]^m$ (which can be approximated arbitrarily well by a discrete measure). It is simple to verify that if $k_0 = k^{1/m}$ is a natural number, then M^* is a uniform distribution on the regular grid of size k_0 in each dimension. It follows that $W_2^2(P^{\text{in}}, M^*) \leq \frac{m}{12k_0^3} = \frac{m}{12k^{3/m}}$. The grid points \mathbf{G} range over the set $S := [1/(2\kappa), 1 - 1/(2\kappa)]^m$. Moreover, Q is a uniform distribution on \mathbf{G} . Now P is the outcome of line (11) so the support of P must lie in the convex hull S of \mathbf{G} . The cost of each unit mass transfer from an atom in the complement of set $[1/(4\kappa), 1 - 1/(4\kappa)]^m$ to one in S is at least $(1/4\kappa)^2$, so $W_2^2(P^{\text{in}}, P) \geq (1/4\kappa)^2 \cdot [1 - (1 - 1/(2\kappa))^m]$. Lastly, note that $(1 - 1/(2\kappa))^m < e^{-m/2\kappa}$. \square

Regularized Rk-means It is possible to extend our approach to accommodate regularization techniques. This can be useful when the data are very high dimensional (Sun et al., 2012; Witten and Tibshirani,

2010). Thus, the clustering formulation can be expressed as a regularized optimal transport problem: $M^* = \arg \min W_2^2(M, P^{\text{in}}) + \Omega(M)$ where the optimization is over the space of discrete measures M that have k support points (the set \mathbf{C} of k centroids), and the regularizer $\Omega(M) \geq 0$ typically decomposes over the m -partition of variables: $\Omega(M) = \sum_{j=1}^m \Omega_j(M_j)$. For instance, $\Omega_j(M_j)$ may be taken to be a multiple of the ℓ_1 norm of M_j 's supporting atoms (e.g., group lasso penalty). The algorithm has the same three steps as before, with some modification in (1') and (3'):

- (1') For each $j \in [m]$, pick M_j to be the (α -approximate) minimizer of $W_2^2(M_j, P_j^{\text{in}}) + \Omega_j(M_j)$, where $\mathbf{C}_j = \text{supp}(M_j)$ is the support of M_j and $|\mathbf{C}_j| = \kappa$ (line 7).
- (3') Finally, return P which is the measure with exactly k support points in \mathbb{R}^d that (γ -approximately) minimizes $W_2^2(P, Q) + \Omega(P)$ (line 11).

Proposition 3.5. *If $\kappa \geq |\text{supp}(M_j^*)| \forall j \in [m]$, then*

$$\frac{W_2^2(P^{\text{in}}, P) + \Omega(P)}{W_2^2(P^{\text{in}}, M^*) + \Omega(M^*)} \leq 2\alpha + 4\gamma + 4\alpha\gamma.$$

If both subproblems for regularized k -means can be solved optimally, our method yields a 10-approximation on the penalized W_2^2 objective. We conclude by noting that our technique extends easily to the W_p^p objective for any $p \geq 1$, but the approximation ratio will be changed according to p .

4 Leveraging relational data

We now explain the ‘‘relational’’ part of the Rk-means algorithm, where we exploit relational structures in the data and the FEQ to achieve significant computational savings. Three classes of relational structures prevalent in RDBMSs are (a) *categorical variables*, (b) *functional dependencies* (FDs), and (c) the topology of the FEQ. We exploit these structures to carefully select the partition $S_1 \cup \dots \cup S_m$ to use for Rk-means, to compute the marginal sub-problems (\mathbf{X}_j, w_j) , the components \mathbf{C}_j of the coresets \mathbf{G} , and the grid weight w_{grid} without materializing the entire coresets \mathbf{G} . When selecting partitions, there are two competing criteria: first, we need a partition so that the approximation ratio α for wkmeans_1 is as small as possible. For example, if $|S_j| = 1 \forall j$, so $m = d$, then we can apply the well-known optimal solution for k -means in 1 dimension using dynamic programming in $O(n^2k)$ time (Wang and Song, 2011); this then provides $\alpha = 1$. On the other hand, we want the remainder of algorithm to be fast by keeping the size of the grid \mathbf{G} , namely $|\mathbf{G}| \leq \kappa^m$, small.

Categorical variables. Real-world relational database queries typically involve many categorical variables (e.g., color, month, or city). In practice, practitioners may endow non-uniform weights for different

categorical variables, or categories (Huang, 1998). In terms of representation, a common way to deal with categorical variables is to one-hot encode them, whereby a categorical feature such as city is represented by an indicator vector $\mathbf{x}_{\text{city}} = [\mathbf{1}_{\text{city}=c_1}, \mathbf{1}_{\text{city}=c_2}, \dots, \mathbf{1}_{\text{city}=c_L}]$ where $\{c_1, \dots, c_L\}$ is the set of cities occurring in the data. The subspace associated with these indicator vectors is known as the *categorical subspace* of a categorical variable. One-hot encoding substantially increases the data matrix size via an increase in the dimensionality of the data. Fortunately, this is not a problem—by treating each categorical variable as a subset of the partition, we can solve the *weighted k-means* subproblem within a categorical subspace efficiently and optimally.

Theorem 4.1. *Given a categorical weighted k-means instance, an optimal solution is to put each of the first $k - 1$ highest weight indicator vectors in its own cluster, and the remaining vectors in the same cluster.*

This means that for a categorical variable with L categories, we can compute the optimal clustering for the sub-problem in only $O(nL \log L)$ time. See Appendix C.1 for more details.

Functional dependencies. Next, we address the second call to `wkmeans2`: its runtime is dependent on the size of the grid \mathbf{G} , which can be up to $O(k^m)$, where m is the number of features from the input. Databases often contain *functional dependencies* (FDs), which we can exploit to reduce the size of \mathbf{G} . An FD is a dimension whose value depends entirely on the value of another dimension. For example, suppose a dataset has such as `storeId`, `zip`, `city`, `state`, and `country`. Here, `storeId` functionally determines `zip`, which determines `city`, which in turns determines `state`, leading to `country`. This common structure is known as an *FD-chain*, and appears often in real-world FEQs. If we were to apply Rk-means naively, these five features would contribute a factor of k^5 to the grid size. However, by using the FD structure of the database, we show that only a factor of $5k$ is contributed to the grid size, because most of the k^5 grid points \mathbf{g} have $w_{\text{grid}}(\mathbf{g}) = 0$ (see (3)). More generally, when there is an *FD-chain* of p features, their overall contribution to the grid size is a factor of $O(kp)$, not $O(k^p)$, and the grid points with non-zero weights can be computed efficiently in time $O(kp)$.

Theorem 4.2. *Suppose all d input features can be partitioned into m FD-chains of size d_1, \dots, d_m , respectively. Then, the number of grid points $\mathbf{g} \in \mathbf{G}$ with non-zero w_{grid} weight is bounded by $\prod_{i=1}^m (1 + d_i(k - 1))$. Furthermore, the set of non-zero weight grid points can be computed in time $\tilde{O}(\prod_{i=1}^m (1 + d_i(k - 1)))$.*

Note that in the above theorem, if there was *no* FD, then d features each form their own chain of size 1, in which case $\prod_{i=1}^m (1 + d_i(k - 1)) = k^m$; thus, the theorem strictly generalizes the no-FD case.

Query structure. Finally, we explain how the FEQ’s structure can be exploited to speed up the computation of subproblems, the grid, and grid weights. In particular, we make use of recent advances in relational query evaluation algorithms (Abo Khamis et al., 2017, 2016; Ngo, 2018; Olteanu and Schleich, 2016). The `InsideOut` algorithm from the `FAQs` framework in particular (Abo Khamis et al., 2017) allows us to compute the grid weights without explicitly the grid points.

For concreteness, we describe the steps of Rk-means as implemented in the database, noting the additional speedups we can get over the description in Alg. 1.

Step 1 (lines 5 and 6). *Project \mathbf{X} into each subspace S_j and compute the weight w of each point.*

In a relational database, the projected sets \mathbf{X}_j already exist in normalized form (Abiteboul et al., 1995). In fact, in sixth normal form (6NF) databases Date et al. (2002), each relation in the database will generally correspond to one variable. So, the sets \mathbf{X}_j and their marginal weights can be computed efficiently. This step perfectly aligns with our strategy of picking the partition $S_1 \cup \dots \cup S_m$ to match the database schema!

Step 2 (line 7). *Find κ centroids in each subspace S_j .*

If the subspace S_j corresponds to a single continuous variable, we can solve the one-dimensional k -means problem quickly and optimally (Wang and Song, 2011); and if S_j corresponds to a categorical feature, then it is solved trivially (and optimally) using Theorem 4.1.

Step 3 (lines 9 and 10). *Construct the coreset \mathbf{G} and the associated weights w_{grid} .*

When constructing \mathbf{G} , it is unnecessary to represent any points in \mathbf{G} that have zero weight. We can use the `InsideOut` database algorithm (Abo Khamis et al., 2016) to efficiently compute nonzero weights, and then extract only those grid points in \mathbf{G} with nonzero weight.

Step 4 (line 11). *Cluster the weighted coreset \mathbf{G} .*

We use a modified version of Lloyd’s weighted k -means that exploits the structure of \mathbf{G} and sparse representation of categorical values to speed up computation. See Appendix C.3 for details.

Runtime analysis. We compare Rk-means to the standard setting of first extracting the matrix \mathbf{X} from the database and then perform clustering on \mathbf{X} directly. The precise runtime statement requires defining a few parameters such as “fractional hypertree width” and “fractional edge cover number” of the FEQ. Details are relegated to Appendix A, and we simply state the main thrust of the runtime result:

Theorem 4.3. *There are classes of feature extraction queries (FEQs) for which the runtime of Rk-means is*




Figure 2: Breakdown of the compute time of Rk-means for each step of the algorithm with $\kappa = k \in \{5, 10, 20, 50\}$. The time to compute \mathbf{X} is provided as reference.

asymptotically less than $|\mathbf{X}|$, and the ratio between $|\mathbf{X}|$ and the runtime of Rk-means can be a polynomial in N , the size of the largest input relation.

The key insight to read from this theorem is that Rk-means can, in principle, run faster than simply exporting the data matrix, without even running *any* clustering algorithm (be it sampling-based, streaming, etc.). Of course, the result only concerns a class of FEQs “on paper”. Section 5 examines real FEQs, which also demonstrate Rk-means’s runtime superiority.

5 Experimental results

We empirically evaluate the performance of Rk-means on three real datasets for three sets of experiments: (1) we break down and analyze the performance of each step in Rk-means; (2) we benchmark the performance and approximation of Rk-means against mlpack (Curtin et al., 2018) (v. 3.1.0), a fast C++ machine learning library; and (3) we evaluate the performance and approximation of Rk-means for setting $\kappa < k$; i.e., different number of clusters for Steps 2 and 4.

The experiments show that the coresets of Rk-means are often significantly smaller than the data matrix. As a result, Rk-means can scale easily to large datasets, and can compute the clusters with a much lower memory footprint than mlpack. When $\kappa = k$, Rk-means is orders-of-magnitude faster than the end-to-end computation for mlpack—up to 115 \times . Typically, the approximation level is very minor. Also, setting $\kappa < k$ can lead to further performance speedups (sometimes exceeding 200 \times !) with only moderate increase in approximation.

Setup. We prototyped Rk-means on top of the LMFAO engine ?. Rk-means is implemented in multi-threaded C++11; this makes mlpack a comparable implementation. All experiments were performed on an AWS x1e.8xlarge instance, which has 1 TiB of RAM and 32 vCPUs. All relations are given sorted by their join attributes.

To construct the data matrix that forms the input to mlpack, we use PostgreSQL 10.6 (psql) to evaluate the FEQ. The seminal k -means++ algorithm (Arthur

	Retailer	Favorita	Yelp
Relations	5	6	6
Attributes	39	15	25
One-hot Enc.	95	1470	1617
# Rows in \mathbf{D}	84M	125M	8.7M
Size of \mathbf{D}	1.5GB	2.5GB	0.2GB
# Rows in \mathbf{X}	84M	127M	22M
Size of \mathbf{X}	18GB	7GB	2.4GB
# Rows in Coreset \mathbf{G}			
$\kappa = 5$	1.43M	14.94K	2.69M
$\kappa = 10$	9.58M	85.88K	11.71M
$\kappa = 20$	38.16M	632.5K	11.89M
$\kappa = 50$	73.75M	7.87M	12.46M

Table 1: Statistics for the input database \mathbf{D} , data matrix \mathbf{X} , and coresets \mathbf{G} for the three dataset.

and Vassilvitskii, 2007) is used for initializing the k -means cluster. We run Rk-means and mlpack + psql five times and report the average approximation and runtime. The timeout for all experiments was set to six hours (21,600 seconds) per trial. Our runtime results omit data loading/saving times. For mlpack + psql, psql must export \mathbf{X} to disk, and then mlpack must then read it from disk! Rk-means has no need to do this, and thus the runtime numbers are skewed in mlpack’s favor. This skew may be significant: loading and saving a large CSV file may take hours.

Datasets. We use three real datasets: (1) *Retailer* is used by a large US retailer for sales forecasting; (2) *Favorita* (Favorita Corp., 2017) is a public dataset for retail forecasting; and (3) *Yelp* is from the public Yelp Dataset Challenge (Yelp, 2017) and used to predict users’ ratings of businesses. Table 1 presents key statistics for the three datasets, including the size of data matrix \mathbf{X} and the coreset \mathbf{G} for each dataset and different κ -values. $|\mathbf{G}|$ is highly data dependent. For *Favorita*, \mathbf{G} is orders-of-magnitude smaller than the data matrix. For *Retailer*, when $\kappa = 20$ and $\kappa = 50$, $|\mathbf{G}|$ approaches $|\mathbf{X}|$, but Rk-means still provides a speedup. Additional dataset details are given in Appendix D.

Breakdown of Rk-means. Figure 2 shows the time it takes Rk-means to cluster the three datasets for different values of k with $\kappa = k$. The total time is broken down into the four steps of the algorithm from Section 4. We provide the time it takes psql to compute \mathbf{X} as ref-

Retailer	k = 5	k = 10	k = 20	k = 50	k=20, $\kappa = 10$	k = 50, $\kappa = 20$
Compute \mathbf{X} (psql)	175.47	175.47	175.47	175.47	175.47	175.47
Clustering (mlpack)	65.41	158.81	385.67	1,453.88	385.67	1,453.88
Rk-means	15.66	54.59	230.17	650.20	63.51	344.31
Relative Speedup	15.38 \times	6.12 \times	2.44 \times	2.51 \times	8.84 \times	4.73 \times
Relative Approx.	0.20	0.08	0.03	0.00	0.03	0.02
Favorita	k = 5	k = 10	k = 20	k = 50	k=20, $\kappa = 10$	k = 50, $\kappa = 20$
Compute \mathbf{X} (psql)	156.86	156.86	156.86	156.86	156.86	156.86
Clustering (mlpack)	1,002.54	6,449.32	11,794.49	>21,600.00	11,794.49	>21,600
Rk-means	27.95	57.72	118.36	334.65	57.65	120.77
Relative Speedup	41.49 \times	114.59 \times	100.98 \times	>64.55 \times	207.30 \times	>178.86 \times
Relative Approx.	2.99	0.35	0.12	–	1.93	–
Yelp	k = 5	k = 10	k = 20	k = 50	k=20, $\kappa = 10$	k = 50, $\kappa = 20$
Compute \mathbf{X} (psql)	33.83	33.83	33.83	33.83	33.83	33.83
Clustering (mlpack)	210.59	640.43	2,107.83	11,474.24	2,107.83	11,474.24
Rk-means	43.37	107.71	195.22	405.11	114.34	241.34
Relative Speedup	5.64 \times	6.26 \times	10.97 \times	28.41 \times	18.73 \times	47.68 \times
Relative Approx.	0.37	0.26	0.13	0.05	0.27	0.20

Table 2: End-to-end runtime and approximation comparison of Rk-means and mlpack on each dataset. The first four columns use different $\kappa = k$ values; the last two show results for setting $\kappa < k$.

erence (gray bar). In many cases, Rk-means can cluster *Retailer* and *Favorita* faster than it takes `psql` to even compute the data matrix! The relative performance of the four steps is data dependent. For *Retailer*, most of the time is spent on constructing \mathbf{G} in Step 3, which is relatively large. For *Favorita*, however, Step 2 takes the longest, as there is one continuous variable with many distinct values, and the DP algorithm for clustering runs in time quadratic in the number of distinct values. Here, performance could be improved by clustering this dimension with a different k -means algorithm; this would increase the approximation somewhat.

Comparison with mlpack. The left columns of Table 2 compares the runtime and approximation of Rk-means against mlpack on the three datasets for different k values with $\kappa = k$. The approximation is given relative to the objective value obtained by mlpack. Speedup is given by comparing the end-to-end performance of Rk-means and mlpack (ignoring disk I/O time), which for mlpack includes the time needed by `psql` to materialize \mathbf{X} . Overall, Rk-means often outperforms even just the clustering step from mlpack, and when end-to-end computation is considered, Rk-means gives up to 115 \times speedup. mlpack timed out after six hours for *Favorita* with $k = 50$. In addition, Rk-means has a much smaller memory footprint than mlpack: for instance, on *Favorita* with $k = 20$, mlpack uses over 900GiB of RAM to cluster the dataset, whereas Rk-means only requires 18GiB. To try and reduce RAM usage, we benchmarked against mlpack with sparse matrices from the Armadillo library (Sanderson and Curtin, 2018); this did reduce RAM usage, but the overhead of working with sparse data structures meant an overall slowdown. Overall, in our simulations, the approximation level is moderate, and consistently well below the 9-approximation bound from Theorem 3.4.

Our simulations show a high level of agreement in the clusterings obtained by the algorithms: for *Retailer* with $k = 20$, the average normalized mutual information is 0.743 between five mlpack k -means clusterings, and 0.711 between Rk-means and mlpack clusterings.

Setting $\kappa < k$ for Step 2. We next evaluate the effect of setting κ to a smaller value than the number of clusters k . This exploits the speed/approximation tradeoff: smaller κ helps reduce the size of \mathbf{G} , at the cost of more approximation. Table 2 presents for each dataset the results for setting $k = 20, \kappa = 10$ and $k = 50, \kappa = 20$, and compares them to the relative performance and approximation over computing k clusters in mlpack.

By setting $\kappa < k$, Rk-means can compute the k clusters up to 208 \times faster than mlpack and 3.6 \times faster than when $\kappa = k$, while the approximation remains moderate. Our results are data dependent—but as the database scales, our speedups will be even more significant.

6 Conclusion

We introduce Rk-means, a method to construct k -means clustering coresets on relational data directly from the database. Rk-means gives a provably good clustering of the entire dataset, without ever materializing the data set; this also yields asymptotic improvements in running time. Experimentally, we observe that the coreset has size up to 180x smaller than the size of the data matrix and this results in orders-of-magnitude improvements in runtime, while still providing empirically good clusterings. Although our work here primarily focuses on k -means clustering, we believe our construction of grid coresets and the accompanying theory is useful for other unsupervised learning tasks and plan to explore such possibilities in future work.

7 Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 682588. B. Moseley was supported in part by NSF grants CCF-1725543, 1733873, 1845146, a Google Research Award, a Bosch junior faculty chair and an Infor faculty award. X. Nguyen was supported in part by grant NSF CAREER DMS-1351362 and a grant from Toyota Research Institute.

Bibliography

- Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison-Wesley, 1995. ISBN 0-201-53771-0. URL <http://webdam.inria.fr/Alice/>.
- Mahmoud Abo Khamis, Hung Q. Ngo, and Atri Rudra. FAQ: questions asked frequently. In *PODS*, pages 13–28, 2016.
- Mahmoud Abo Khamis, Hung Q. Ngo, and Atri Rudra. Juggling functions inside a database. *SIGMOD Rec.*, 46(1):6–13, 2017.
- Mahmoud Abo Khamis, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. Ac/dc: In-database learning thunderstruck. In *2nd Workshop on Data Mgt for End-To-End ML, DEEM’18*, pages 8:1–8:10, 2018a.
- Mahmoud Abo Khamis, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. In-database learning with sparse tensors. In *PODS*, pages 325–340, 2018b.
- Sara Ahmadian, Ashkan Norouzi-Fard, Ola Svensson, and Justin Ward. Better guarantees for k-means and euclidean k-median by primal-dual algorithms. In *FOCS*, pages 61–72, 2017.
- D. Arthur and S. Vassilvitskii. k-means++: The advantages of careful seeding. In *SODA*, page 1027–1035, 2007.
- Olivier Bachem, Mario Lucic, and Andreas Krause. Scalable k-means clustering via lightweight coresets. In *SIGKDD*, pages 1119–1127, 2018.
- Bahman Bahmani, Benjamin Moseley, Andrea Vattani, Ravi Kumar, and Sergei Vassilvitskii. Scalable k-means++. *PVLDB*, 5(7):622–633, 2012.
- Maria-Florina Balcan, Steven Ehrlich, and Yingyu Liang. Distributed k-means and k-median clustering on general communication topologies. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 1995–2003, 2013.
- Vladimir Braverman, Gereon Frahling, Harry Lang, Christian Sohler, and Lin F. Yang. Clustering high dimensional dynamic data streams. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 576–585, 2017.
- Field Cady. *The Data Science Handbook*. John Wiley & Sons, 2017.
- Ryan R. Curtin, Marcus Edel, Mikhail Lozhnikov, Yanis Mentekidis, Sumedh Ghaisas, and Shangdong Zhang. mlpack 3: a fast, flexible machine learning library. *J. Open Source Software*, 3:726, 2018.
- C.J. Date, H. Darwen, and N. Lorentzos. *Temporal Data & The Relational Model*. Elsevier, 2002.
- E. del Barrio, J. A. Cuesta-Albertos, C. Matran, and A. Mayo-Isacar. Robust clustering tools based on optimal transportation. *Statistics and Computing*, pages 1–22, 2017.
- Alina Ene, Sungjin Im, and Benjamin Moseley. Fast clustering using mapreduce. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Diego, CA, USA, August 21-24, 2011*, pages 681–689, 2011.
- B. Everitt, S. Landau, M. Leese, and D. Stahl. *Cluster Analysis*. Wiley & Sons, 2011.
- Favorita Corp. Corporacion Favorita Grocery Sales Forecasting: Can you accurately predict sales for a large grocery chain? <https://www.kaggle.com/c/favorita-grocery-sales-forecasting/>, 2017.
- S. Graf and H. Luschgy. *Foundations of quantization for probability distributions*. Springer-Verlag, New York, 2000.
- Martin Grohe and Dániel Marx. Constraint solving via fractional edge covers. *ACM Trans. Algorithms*, 11(1):4:1–4:20, 2014.
- Sudipto Guha, Adam Meyerson, Nina Mishra, Rajeev Motwani, and Liadan O’Callaghan. Clustering data streams: Theory and practice. *IEEE Trans. Knowl. Data Eng.*, 15(3):515–528, 2003.
- S. Har-Peled and S. Mazumdar. On coresets for k-means and k-median clustering. In *SODA*, 2004.
- Sariel Har-Peled and Soham Mazumdar. Coresets for k-means and k-median clustering and their applications. *CoRR*, abs/1810.12826, 2018. URL <http://arxiv.org/abs/1810.12826>.
- J. A. Hartigan. *Clustering algorithms*. Wiley, New York, 1975.
- N. Ho, X. Nguyen, M. Yurochkin, H. H. Bui, V. Huynh, and D. Phung. Multilevel clustering via Wasserstein means. In *ICML*, pages 1501–1509, 2017.

- Z. Huang. Extensions to the k-means algorithm for clustering large data sets of categorical values. *Data mining and Knowledge discovery*, 2:283–304, 1998.
- L. Kaufman and P. J. Roussew. *Finding Groups in Data - An Introduction to Cluster Analysis*. Wiley & Sons, 1990.
- Mahmoud Abo Khamis, Ryan Curtin, Benjamin Moseley, Hung Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. On functional aggregate queries with additive inequalities. In *PODS*, 2019.
- Daphne Koller and Nir Friedman. *Probabilistic graphical models*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, 2009. Principles and techniques.
- Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Trans. Inf. Theory*, 28(2):129–136, 1982.
- Dániel Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM*, 60(6):42:1–42:51, 2013.
- Adam Meyerson, Liadan O’Callaghan, and Serge A. Plotkin. A k -median algorithm with running time independent of data size. *Machine Learning*, 56(1-3): 61–87, 2004.
- Hung Q. Ngo. Worst-case optimal join algorithms: Techniques, results, and open problems. In *PODS*, pages 111–124, 2018.
- Hung Q. Ngo, Ely Porat, Christopher Ré, and Atri Rudra. Worst-case optimal join algorithms. *Journal of the ACM (JACM)*, 65(3):1–40, 2018.
- Dan Olteanu and Maximilian Schleich. Factorized databases. *SIGMOD Rec.*, 45(2):5–16, 2016.
- Carlos Ordonez. Integrating k-means clustering with a relational DBMS using SQL. *IEEE Trans. Knowl. Data Eng.*, 18(2):188–201, 2006.
- Carlos Ordonez and Edward Omiecinski. Efficient disk-based k-means clustering for relational databases. *IEEE Trans. Knowl. Data Eng.*, 16(8):909–921, 2004.
- D. Pollard. Quantization and the method of k-means. *IEEE Trans. Inf. Theory*, 28(2):199–204, 1982.
- C. Sanderson and R.R. Curtin. A user-friendly hybrid sparse matrix class in C++. In *Proceedings of the 2018 International Congress on Mathematical Software (ICMS)*, pages 422–430. Springer, 2018. ISBN 978-3-319-96418-8.
- Maximilian Schleich, Dan Olteanu, and Radu Ciucanu. Learning linear regression models over factorized joins. In *SIGMOD*, pages 3–18, 2016.
- Maximilian Schleich, Dan Olteanu, Mahmoud Abo Khamis, Hung Q. Ngo, and XuanLong Nguyen. A layered aggregate engine for analytics workloads. In *SIGMOD*, pages 1642–1659, 2019.
- Christian Sohler and David P. Woodruff. Strong coresets for k-median and subspace approximation: Goodbye dimension. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 802–813, 2018.
- W. Sun, J. Wang, and Y. Fang. Regularized k-means clustering of high-dimensional data and its asymptotic consistency. *Electronic Journal of Statistics*, 9: 148–167, 2012.
- Mikkel Thorup. Quick k-median, k-center, and facility location for sparse graphs. In *Automata, Languages and Programming, 28th International Colloquium, ICALP 2001, Crete, Greece, July 8-12, 2001, Proceedings*, pages 249–260, 2001.
- Todd L. Veldhuizen. Leapfrog Triejoin: A simple, worst-case optimal join algorithm. *arXiv preprint arXiv:1210.0481*, 2012.
- Cédric Villani. *Optimal transport*, volume 338 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2009.
- Haizhou Wang and Mingzhou Song. Ckmeans.1d.dp: Optimal k-means clustering in one dimension by dynamic programming. *The R Journal*, 3(2):29, 2011.
- D. Witten and R. Tibshirani. A framework for feature selection in clustering. *Journal of the American Statistical Association*, 105:713–726, 2010.
- Xindong Wu, Vipin Kumar, J. Ross Quinlan, Joydeep Ghosh, Qiang Yang, Hiroshi Motoda, Geoffrey J. McLachlan, Angus F. M. Ng, Bing Liu, Philip S. Yu, Zhi-Hua Zhou, Michael Steinbach, David J. Hand, and Dan Steinberg. Top 10 algorithms in data mining. *Knowl. Inf. Syst.*, 14(1):1–37, 2008.
- J. Ye, P. Wu, J. Wang, and J. Li. Fast discrete distribution clustering using barycenter with sparse support. *IEEE Trans. Signal Proc.*, 65(9):2317–2332, 2017.
- Yelp. Yelp dataset challenge, <https://www.yelp.com/dataset/challenge/>, 2017.