Coresets over Multiple Tables for Feature-rich and Data-efficient Machine Learning

Jiayi Wang
Tsinghua University
jiayi-wa20@mails.tsinghua.edu.cn

Chengliang Chai
Tsinghua University
ccl@tsinghua.edu.cn

Nan Tang
QCRI
ntang@hbku.edu.qa

Jiabin Liu
Tsinghua University
liujb19@mails.tsinghua.edu.cn

Guoliang Li
Tsinghua University
liguoliang@tsinghua.edu.cn

ABSTRACT

Successful machine learning (ML) needs to learn from good data. However, one common issue about train data for ML practitioners is the lack of good features. To mitigate this problem, feature augmentation is often employed by joining with (or enriching features from) multiple tables, so as to become feature-rich ML. A consequent problem is that the enriched train data may contain too many tuples, especially if the feature augmentation is obtained through 1 (or many)-to-many or fuzzy joins. Training an ML model with a very large train dataset is data-inefficient. Coreset is often used to achieve data-efficient ML training, which selects a small subset of train data that can theoretically and practically perform similarly as using the full dataset. However, coreset selection over a large train dataset is also known to be time-consuming.

In this paper, we aim at achieving both feature-rich ML through feature augmentation and data-efficient ML through coreset selection. In order to avoid time-consuming coreset selection over a feature augmented (or fully materialized) table, we propose to efficiently select the coreset without materializing the augmented table. Note that coreset selection typically uses weighted gradients of the subset to approximate the full gradient of the entire train dataset. Our key idea is that the gradient computation for coreset selection of the augmented table can be pushed down to partial feature similarity of tuples within each individual table, without join materialization. These partial feature similarity values can be aggregated to estimate the gradient of the augmented table, which is upper bounded with provable theoretical guarantees. Extensive experiments show that our method can improve the efficiency by nearly 2 orders of magnitudes, while keeping almost the same accuracy as training with the fully augmented train data.

(a) Base table **Movie** and other four tables, **Direct, Person, Production and Company**, that can be used to augment the feature of **Movie**

(b) Feature augmented table **Movie⁺**

Figure 1: Feature-rich ML through feature augmentation.

1 INTRODUCTION

Feature-rich machine learning (ML) [14, 15, 29, 35] means that ML models are trained with enough and good features. Given train data, *data-efficient ML* [3, 4, 39] aims at training ML models faster without sacrificing the model accuracy. Putting them together, the goal is to efficiently train robust ML models.

For achieving **feature-rich** ML, the widely used practice is to enrich features by joining a base table with multiple tables, *a.k.a. feature augmentation* [12, 15, 30].

EXAMPLE 1. *[Feature-rich ML through Feature Augmentation.] Consider an ML task that predicts the* Score *value of a movie based on attributes (i.e., features)* MovieID, Title, Length *and* BoxOffice, *as shown in the* Movie *table in Figure 1(a). Intuitively, because many important features are missing, such as the features of directors and the actors of a movie, it is hard to train a good ML model.*

Consider four tables Direct, Person, Production *and* Company, *which can be joined, directly or indirectly, with the* Movie *table through predefined joins, as shown in Figure 1(a). The primary key of each table is annotated with an underline (e.g.,* MovieID *for table* Movie*). These tables can be joined with either 1-to-1 relationship or 1-to-many (i.e., 1-to-∗) relationship. The feature-rich table with new features augmented through joins, denoted by* Movie⁺*, is shown in Figure 1(b).*

For **data-efficient** ML, besides traditional methods (*e.g.,* stochastic gradient descent and its variants), there are many recent efforts on selecting a train data subset that can theoretically and practically perform on par with the full dataset, *a.k.a. coresets* [18, 41].

Figure 2: Various design choices.

Figure 3: Comparison of various design choices.

In order to achieve both **feature-rich** and **data-efficient** ML, an intuitive solution is to first conduct feature augmentation through joins across multiple relational tables, followed by performing coreset selection over the train data with enriched features. In Figure 2, this strategy is depicted by first following the "$\Rightarrow$" arrow from the base table T and then the "$\rightarrow$" arrow.

Before we discuss the problem we study in this paper, let's analyze the benefits and limitations of the aforementioned strategies, using real experiments (see more details in § 6).

EXAMPLE 2. *[Benefits and Limitations of Existing Strategies.] Please refer to Figure 3 for the following discussions.*

Train with the Base Table. *If we train an ML model using the base table T for a multi-classification task, we can achieve an accuracy 0.61 (Figure 3(a)–❶) and use 11 minutes for training (Figure 3(b)–❶). We consider this as a baseline.*

Train with the Coreset. *If we first compute the coreset $C(T)$ of the base table T, and then train an ML model using the computed coreset, we can also achieve an accuracy 0.61 (Figure 3(a)–❷) but use 2 minutes in total for selecting the coreset and training with the coreset (Figure 3(b)–❷). This shows that using coreset can significantly reduce the training time without sacrificing the accuracy.*

Train with the Augmented Table. *As we know, the process of feature augmentation consists of 1 (many)-to-many joins [14, 15, 35, 49]. In this situation, the size of the augmented table T^+ is likely to be much larger than the base table T. Although training with the augmented table can achieve a much higher accuracy 0.68 (Figure 3(a)–❸), it takes around 2.8 days for training over 10 million tuples (Figure 3(b)–❸).*

Train with the Coreset of the Augmented Table. *If we first compute the coreset of the augmented table, and then train with this coreset, we can achieve the same accuracy 0.68 (Figure 3(a)–❹). However, because coreset computation on a large table is time-consuming which takes 2.2 days in this case, putting the time of feature augmentation (5 minutes) and training (0.5 hours) together, it takes around 2.2 days (Figure 3(b)–❹).*

Example 2 tells us that feature augmentation can significantly increase the accuracy and training with a coreset can significantly reduce the time. However, computing the coreset over a large table (*e.g.*, the augmented table) is time-consuming.

In order to efficiently compute the coreset of the augmented table, the **problem** we aim to tackle is whether we can **compute the coreset of the augmented table without join materialization.**

This strategy is depicted in Figure 2 by following the "$\rightsquigarrow$" arrow from the base table.

Key Idea. Our solution is inspired by the classical SQL query optimization technique **pushdown** that moves predicates in the WHERE clause closer to the tables they refer to. Next let's build the connection between pushdown and our problem. Generally speaking, coreset computation is to select a subset of tuples, and use the weighted gradients of these tuples to approximate the full gradient of the entire train data. In our context, **pushdown for coreset over multiple tables** means that we can approximately compute the gradient of each individual table and sum up these gradients from multiple tables to compute the full gradient of the joined table. That is, the *gradient computation* is pushed down.

Challenges. Applying *pushdown* to compute the coreset of the augmented table without join materialization faces two challenges.

(C1) How to deal with each individual table so that the estimated full gradients can be bounded with theoretical guarantees, and thus the accuracy will not be sacrificed?

(C2) How to efficiently aggregate information distributed in different tables to well approximate the full gradient?

Our Methodology. To address the above challenges, we propose feature-Rich and data-Efficient machine learning with Coreset selection withOut join materializatioN, namely RECON, which groups the tuples in each table based on predefined group key and compute the gradient bounds of these groups based on tuple-wise partial feature similarity. We prove that the full gradient can be bounded using these partial feature similarity values of different groups (for **C1**). Based on the computed partial feature similarity values, we prove that the coreset computation problem over multiple tables is NP-hard and has the submodular property, so we use a greedy algorithm with an approximate ratio to aggregate the partial feature similarity values in different tables, and then compute a well-performed coreset (for **C2**).

Figure 3 shows that RECON can significantly reduce the total time from more than 2 days to 1.8 hours (Figure 3(b)–❺) without sacrificing the accuracy (Figure 3(a)–❺). See § 6 for more details.

Contributions. In this paper, we make a first attempt to study coreset selection over multiple tables without full materialization. In sum, we make the following notable contributions.

(1) We introduce coreset selection over a single table, its sample-based optimization, and our key idea of supporting coreset selection over multiple tables without full materialization in § 3.

(2) We provide theoretical guarantees on the gradient approximation of the augmented table using partial feature similarity values computed from each single table. We also prove that the coreset computation over multiple tables is an NP-hard problem with the submodular property. Putting them together, we theoretically show that gradient approximation error of coreset for the augmented table can be upper bounded by the computation of each individual table, without physically joining them in § 4.

(3) We propose an efficient greedy algorithm to compute the coreset of the augmented table without materializing the joins. To be specific, we utilize a dynamic programming algorithm to efficiently aggregate partial feature similarity values from multiple tables and finally bound the full gradient in § 5.

(4) We conducted extensive experiments on 5 real world datasets to evaluate the efficiency of our proposed method. The experimental results demonstrate that our method can improve the efficiency to nearly 2 orders of magnitudes by training over the selected coreset ($\sim$ 0.1% of the entire train data), while keeping almost the same accuracy as training with the fully augmented train data in § 6.

2 RELATED WORK

Generally speaking, there are two ways of table enrichment for better training ML models, either by adding more columns (*i.e.,* feature augmentation) or more tuples (*i.e.,* data acquisition [48]). Both are common in practice and they are complementary to each other. Our proposal falls into the category of feature augmentation.

Feature augmentation. A ***brute force*** solution is to execute joins to augment new features. Another line of research focuses on ***avoiding unnecessary joins*** (*e.g.,* Kumar et al. [30] and Shah et al. [50]), when the foreign join key has already contained all the information of the external table. There are also studies [15, 35] on ***iterative feature augmentation***, which iteratively select an optimal subset of tables, such that features augmented from these tables can significantly improve the model performance.

Different from (iterative) feature augmentation, by following the setting in [21, 29, 34], we assume that which joins are useful (*i.e.,* which attributes should be added) are given. Besides, iterative feature augmentation [15, 35] needs to join tables, train over the result and test the performance iteratively, which is rather time-consuming. Therefore, RECON can be leveraged to accelerate this process in each iteration. In terms of avoiding unnecessary joins, we can take it as a filtering of our method. That is, the outputs of the method are the input of our method. Note that [30, 50] do not support to avoid one-to-many, many-to-many and fuzzy joins. In summary, the above methods are complementary to our proposal.

Note that when the number of features is large, **feature selection** (*a.k.a.* variable selection or attribute selection), is the process of selecting a subset of *good* features for use in model construction and has been extensively studied in the ML community [20], which is orthogonal to feature augmentation tackled in this paper.

Coreset selection. Existing coresete selection algorithms are designed for one table. Huang et al. [23] proposed to select and update the coreset while training. The goal is to use the loss of training tuples in the coreset to approximate the overall training loss of the entire dataset. Since they have to train the model, it is rather

time-consuming. To address this, works [8, 9] focused on selecting the coreset without training in advance, but they are customized to particular model types respectively. The high level idea is to compute a sensitive score for each tuple. The higher the score, the more likely the tuple should be a member of the coreset. Dataset condensation [56, 57] tries to synthesize a small set of train data, instead of selecting a small set of train data. Hence, the related work is closer to knowledge distillation than coreset selection. Moreover, these papers are only tested on image data. In terms of tabular data, it is not verified whether it can synthesize a small set of train data while well preserving the labels. The other typical line of works [27, 39, 40] focused on selecting the coreset to approximate the full gradient, which is modeled as an optimization problem that can be solved by a framework with three nested loops (see § 3.2).

None of them considers coreset selection over multiple tables. Different from them, we make the first attempt to select coresets over multiple tables without full materialization.

Factorized ML (FML) achieves efficient ML training by decoupling the ML computations through joins to the base tables [14, 25, 28, 29, 34, 43, 47, 49]. The key idea is to reduce redundant linear algebra computations during training over the multi-table joins. Most of these methods only focus on specific ML models or platforms, *e.g.,* Olteanu et al. [43, 49] focus on linear regression models, and [28, 47] are for in-memory databases. Recently, Kumar et al. [14, 25, 29, 34] build a general FML framework by decoupling linear algebra computations from various ML algorithms.

The difference from us is that they still need to *train over the full join results*. Moreover, we aim at accelerating ML training by reducing the amount of train data through selecting the coreset. However, they focus on the batch gradient descent algorithm for training rather than supporting the stochastic gradient descent., which is widely used in practice due to its high efficiency.

In addition, we have also empirically verified that our proposal outperforms FML-based methods for the batch gradient descent scenario (see § 6 for more details).

Data management techniques for ML. Besides existing works [6, 13, 14, 17, 29, 31, 54] that leverage database techniques to improve the efficiency, there are also many works that aim at improving the effectiveness of ML models, including data discovery [15, 36], data labeling [11, 33], data compression [52, 55], data cleaning [10] and data exploration [44, 45].

3 CORESET SELECTION FRAMEWORK
3.1 Gradient Descent for Machine Learning

Gradient descent is by far the most popular optimization strategy used in machine learning. Generally speaking, based on a convex and differentiable function, it iteratively tweaks the parameters to minimize a given function to its local minimum.

Let $T = \{t_1, t_2, ..., t_n\}$ be a set of labeled train tuples, where $t_i = (\mathbf{x}_i, \mathbf{y}_i)$, $\mathbf{x}_i \in \mathbb{R}^d$ denotes the feature vector and $\mathbf{y}_i$ is the corresponding label. The objective of training on T is to compute the best parameter θ^* *w.r.t.* an ML model so as to minimize the loss:

$$\theta^* = \arg \min_{\theta \in \vartheta} l(\theta), l(\theta) = \frac{1}{n} \sum_{i=1}^{n} l_i(\theta, t_i) \tag{1}$$

Algorithm 1: Basic Coreset Algorithm of One Table

Input: The train data T, coreset size K.
Output: A coreset $C \subseteq T$, weight $W = \{w_j\}, |C| = |W| = K$.

1 $C = \emptyset$;
2 **while** $|C| < K$ **do**
3   /*1st loop*/
4   **for** *each tuple* $t \in T \setminus C$ **do**
5   /*2nd loop*/
6   Compute $U(t|C)$ considering all tuples in T; /*3rd loop*/
7      $t^* = \arg\max_{t \in T \setminus C} U(t|C)$;
8      $C = C \cup \{t^*\}$;
9 **for** $j = 1$ to $|C|$ **do**
10      $w_j = \sum_{i=1}^{n} \mathbb{I}[j = \arg\min_{c_{j'} \in C} \max_{\theta \in \vartheta} \|\nabla l_i(\theta) - \nabla l_{\gamma(j')}(\theta)\|]$;
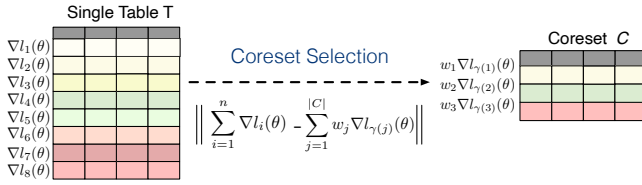11 **return** C, W;

Single Table T   Coreset Selection   Coreset C

$\nabla l_1(\theta)$
$\nabla l_2(\theta)$
$\nabla l_3(\theta)$
$\nabla l_4(\theta)$
$\nabla l_5(\theta)$
$\nabla l_6(\theta)$
$\nabla l_7(\theta)$
$\nabla l_8(\theta)$

$\left\| \sum_{i=1}^{n} \nabla l_i(\theta) - \sum_{j=1}^{|C|} w_j \nabla l_{\gamma(j)}(\theta) \right\|$

$w_1 \nabla l_{\gamma(1)}(\theta)$
$w_2 \nabla l_{\gamma(2)}(\theta)$
$w_3 \nabla l_{\gamma(3)}(\theta)$

Figure 4: Example of coreset selection for a single table.

where ϑ denotes the parameter space. For ease of representation, we just use $l_i(\theta)$ to denote the loss of the i-th train example, *i.e.*, $l_i(\theta, t_i)$. Typically, the gradient descent method is always applied to optimize Eq. 1, where the **full gradient**, denoted by $\nabla l(\theta)$, is required to be computed iteratively.

Although some classic incremental gradient methods such as stochastic gradient descent (SGD), can be utilized to accelerate this process, it is still expensive when there are massive train tuples.

3.2 Coreset of One Table T

Coreset. The main problem of learning using a large train dataset T is low efficiency. Hence, instead of learning from entire T, one research direction seeks to answer the question that whether we can compute a small subset $C(T)$ of T such that learning using $C(T)$ can hopefully have the same performance as learning using T. This small subset is called the **coreset** [18, 41]. In the rest of the paper, we will simply write $C(T)$ as C, when it is clear from the context.

To compute the coreset, the SOTA solutions are mainly based on gradient approximation [26, 39]. Intuitively, if θ is the parameter of an ML model trained using the full dataset, and θ' is the parameter of the same ML model trained using the subset (or coreset), the goal is $\nabla l(\theta) = \nabla l(\theta')$. Based on gradient approximation, existing solutions can lead to good performance with theoretical guarantees, *i.e.*, $\nabla l(\theta)$ is upper-bounded by $\nabla l(\theta')$. Next let's formally define it.

Coreset selection based on gradient descent. Let $\nabla l(\theta) = \sum_{i=1}^{n} \nabla l_i(\theta)$ be the full gradient training using the entire training dataset, the problem of coreset selection is to minimize the **gradient approximation error** [39] between the full gradient *w.r.t.* T and the weighted sum of gradients *w.r.t.* the coreset C (or coreset gradient).

$$C^* = \underset{C \subseteq T, w_j \geq 0}{\arg\min} \underbrace{\max_{\theta \in \vartheta} \| \underbrace{\sum_{i=1}^{n} \nabla l_i(\theta)}_{\text{full gradient}} - \underbrace{\sum_{j=1}^{|C|} w_j \nabla l_{\gamma(j)}(\theta)}_{\text{coreset gradient}} \|}_{\text{gradient approximation error}}, \quad (2)$$

$$s.t. \ |C| \leq K$$

Eq. 2 tries to minimize the gradient approximation error using a coreset of size at most K by considering all possible parameters $\theta \in \vartheta$ (*i.e.*, $\max_{\theta \in \vartheta}$), where "$\| \cdot \|$" denotes the normed difference.

The **full gradient** has been defined earlier as $\nabla l(\theta) = \sum_{i=1}^{n} \nabla l_i(\theta)$. Next, let's focus on explaining how to compute the **coreset gradient** in Eq. 2. We use $\gamma(j) = i, j \in [1, |C|], i \in [1, n]$ to denote that the j-th tuple in C (denoted by c_j) is the i-th tuple in T, *i.e.*, t_i. In other words, γ is an index mapping from C to T. Besides, Eq. 2 potentially contains another important mapping ϕ similar to γ, *i.e.*, $\phi(i) = j, i \in [1, n], j \in [1, |C|]$, which has a close relationship with the weight w. To be specific, let $\phi(i) = j$ denote that we will assign t_i to c_j and use $\nabla l_{\gamma(j)}$ to represent ∇l_i. Each t_i will be assigned to one and only one c_j, but each c_j might be assigned with multiple tuples in T. Based on ϕ, w_j is defined as the weight of the c_j, which is the number of tuples in T assigned to the c_j, *i.e.*, $w_j = |\{t_i | \phi(i) = j, i \in [1, n]\}|$.

Next let's use an example to better illustrate Eq. 2.

EXAMPLE 3. *Let us consider a special case of the gradients of each tuple, as shown in Figure 4. Suppose that for any* θ, $\nabla l_1(\theta) \approx \nabla l_2(\theta) \approx \nabla l_3(\theta)$, $\nabla l_4(\theta) \approx \nabla l_5(\theta)$ *and* $\nabla l_6(\theta) \approx \nabla l_7(\theta) \approx \nabla l_8(\theta)$. *In this situation, if we want to find an optimal coreset with a size of 3, i.e.,* $K = 3$ *based on Eq. 2, the solution can be* $C^* = \{t_1, t_4, t_6\}$ *(*$\gamma(1) = 1, \gamma(2) = 4$ *and* $\gamma(3) = 6$*), associated with* $w_1 = 3, w_2 = 2, w_3 = 3$ *because* $\phi(1) = \phi(2) = \phi(3) = 1, \phi(4) = \phi(5) = 2$ *and* $\phi(6) = \phi(7) = \phi(8) = 3$. *In this way,* C^* *will be one of the optimal coresets that can well approximate the full gradient because* $\| \sum_{i=1}^{8} \nabla l_i(\theta) - \sum_{j=1}^{3} w_j \nabla l_{\gamma(j)}(\theta) \| \approx 0$, *which is minimized.*

We can observe from Example 3 that, if $\phi(i) = j$, ∇l_i and $\nabla l_{\gamma(j)}$ are likely to be close such that the gradient approximation error, *i.e.*, Eq. 2, tends to be minimized. Intuitively, computing the coreset is similar to computing the K exemplars [46] of the gradients, if all the gradients of tuples can be computed.

For training with the popular SGD method, the coreset C is first randomly shuffled. Then during each step of gradient descent, suppose that we need to use $c_j \in C$ for gradient update. We first compute the gradient of c_j, say ∇l. Then we use $w_j \nabla l$ to update the parameters of the ML model. The above process is repeated until the ML model converges.

Basic Coreset Algorithm of One Table. The basic coreset algorithm is illustrated in Figure 5(a) and Algorithm 1. Initialized with an empty coreset C, a coreset of size K is achieved using three nested for-loops.

- **1st for-loop (lines 2-8).** Each iteration of the 1st for-loop will add the tuple with the maximum "utility" to the coreset (lines 7-8). The "utility" of a tuple t denotes the reduction
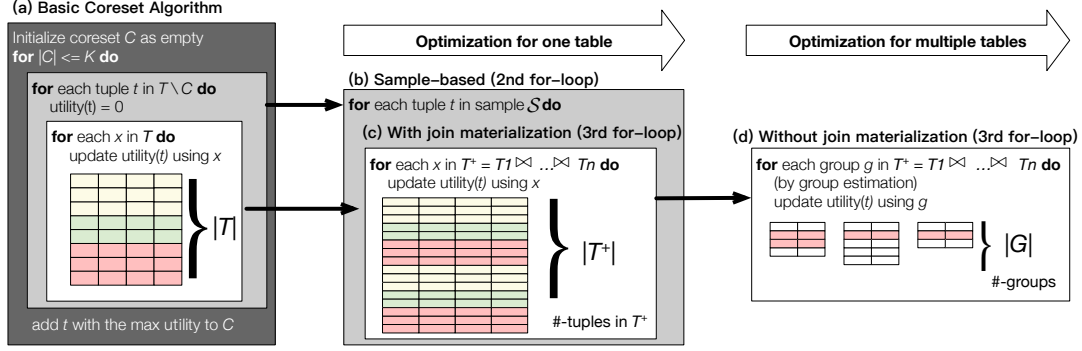
Figure 5: Basic coreset selection (a), typical optimization techniques (b, c), and our optimization for multiple tables (d).

of gradient approximation error in Eq. 2 after adding t into the coreset C, denoted by *i.e.,* $U(t|C)$.

- **2nd for-loop (lines 4-6).** Each iteration of the 2nd-iteration will compute the utility of a tuple t that is not in coreset C.
- **3rd for-loop (line 6).** It iterates all tuples in T to compute the utility of tuple t used in the 2nd for-loop.
- **Weights computation (lines 9-10).** It computes the weight of each tuple in the coreset, which will be used to approximate the full gradient.

Apparently, the solution with 3 loops is rather time-consuming. Fortunately, coresets satisfy the submodular property [24] (Theorem 2 in § 4), based on which an efficient method can accelerate the 2nd loop that uniformly samples tuple set $\mathcal{S}$ from $T \setminus C$ and selects the best one from the sampled ones [38]. It holds a $(1 - \frac{1}{e} - \epsilon)$ approximate ratio, where ϵ is related to the sampling ratio.

Coreset by sampling. The optimization of sampling-based coreset selection algorithm is illustrated in Figure 5(b), where the sample $\mathcal{S}$ is a subset of $T \setminus C$ ("\" denotes the operation set difference).

3.3 Coreset of Multiple Tables

As discussed above, to achieve feature-rich ML, the base table has to be augmented to get useful features through joining other tables.

With materialization. A natural solution is first to do feature augmentation by executing the joins, and then use the above coreset selection on the materialized result. Figure 5(c) depicts this solution.

Note that, for feature augmentation, there might have one-to-one, one-to-many, many-to-many and fuzzy joins. Hence, the materialized view might be very large. Consequently, the efficiency of coreset selection is low (see Figure 3(b)-❹).

Without materialization. Our key idea to improve the efficiency of coreset of multiple tables is to estimate the utility of each "group" without join materialization, where a group refers to a set of tuples in the joined results having the same attribute values on a predefined set of attributes, as shown in Figure 5(d). Conceptually, the utility can be estimated by first computing the feature similarity of tuples in each individual table, and then aggregating them using a dynamic programming algorithm without join materialization. By doing so, we can significantly reduce the computation in the 3rd for-loop, thus improving the overall efficiency (see Figure 3(b)-❺).

4 GRADIENT APPROXIMATION ERROR BOUNDED BY GROUPS

In this section, we will build the theoretical bound of gradient computation for coresets over multiple tables. Afterwards, we will introduce the algorithms by following the theoretical bounds (§ 5).

Let T be the base table, $T_1, T_2, \ldots, T_m$ be the tables that can be used to augment the features of T, and T^+ be the feature-augmented table through predefined joins, with $|T^+| = N$. We will prove that the gradient approximation error of a coreset *w.r.t.* T^+ can be upper-bounded using the groups (or partitions) of T^+.

Recall that minimizing Eq. 2 is closely related to the parameter θ. Unfortunately, the entire parameter space ϑ is too expensive to explore. We will first prove that the gradient approximation error is upper-bounded for a fixed parameter θ and given groups (§ 4.1). We will then generalize the above result to the parameter space ϑ for given groups (§ 4.2). We will close this section by discussing the connection between groups in the augmented single table (*i.e.,* T^+) and multiple tables (§ 4.3).

Building upon the above results, later § 5 will describe how to compute groups from multiple tables (*i.e.,* $T, T_1, \ldots, T_m$) and use them to bound the gradient approximation error of the coreset of T^+ without materializing T^+.

<u>Notation.</u> Similar to § 3, we use ∇l_i^+ to denote the loss of the i-th training example in T^+, *i.e.,* $l_i^+(\theta, t_i)$.

4.1 Upper Bound of a Fixed θ and Given Groups

Recap that we have a one-to-one mapping $\phi_\theta : T^+ \to C$ between tuples in T^+ and C. $\phi_\theta(i) = j$ denotes that t_i^+ in T^+ should be assigned to the j-th tuple in the coreset. With this mapping, we get:

$$\sum_{j=1}^{|C|} w_j \nabla l_{\gamma(j)}^+ (\theta) = \sum_{i=1}^{N} \nabla l_{\gamma(\phi_\theta(i))}^+ (\theta)$$

so Eq. 2 can be rewritten as [39]:

$$\|\sum_{i=1}^{N} \nabla l_i^+(\theta) - \sum_{j=1}^{|C|} w_j \nabla l_{\gamma(j)}^+(\theta)\| = \|\sum_{i=1}^{N} \left(\nabla l_i^+(\theta) - \nabla l_{\gamma(\phi_\theta(i))}^+(\theta) \right)\| \tag{3}$$

Let $\mathcal{A}$ be the predefined attribute set (or the grouping key), based on which tuples in T^+ are divided into a set $\mathcal{G}$ of disjoint groups $g = |\mathcal{G}|$ such that each group $\mathcal{G}_i \in \mathcal{G}$ contains tuples with the

same values on $\mathcal{A}$ (see § 6.1 for more implementation details). We then use $G_i, i \in [1,g]$ to denote the set of indexes (corresponding to T^+) of tuples in $\mathcal{G}_i$, i.e., $G_i = \{k | t_k^+ \in T^+, t_k^+ \in \mathcal{G}_i\}$. That is, $\cup_{i=1}^g G_i = \{1, 2, \ldots, N\}$. With such grouping, different from [39], the summation from 1 to N in Eq. 3 can be rewritten as the sum of g summations over the computation results in each $\mathcal{G}_i$:

$$
\begin{aligned}
&\|\sum_{i=1}^{N} \left(\nabla l_i^+(\theta) - \nabla l_{\gamma(\phi_\theta(i))}^+(\theta) \right)\| \\
&= \|\sum_{i=1}^{g} \sum_{k \in G_i} \left(\nabla l_k^+(\theta) - \nabla l_{\gamma(\phi_\theta(k))}^+(\theta) \right)\| \\
&\leq \sum_{i=1}^{g} \| \sum_{k \in G_i} \left(\nabla l_k^+(\theta) - \nabla l_{\gamma(\phi_\theta(k))}^+(\theta) \right)\| \\
&\leq \sum_{i=1}^{g} |G_i| \max_{k \in G_i} \|\nabla l_k^+(\theta) - \nabla l_{\gamma(\phi_\theta(k))}^+(\theta)\|
\end{aligned}
\tag{4}
$$

Eq. 4 comes from the triangle equation. More specifically, $\|\nabla l_k^+(\theta) - \nabla l_{\gamma(\phi_\theta(k))}^+(\theta)\|$ denotes the gradient difference between a tuple $t_k^+ \in T^+$ and the tuple $t_{\gamma(\phi_\theta(k))}^+$ that t_k^+ assigns to. Hence, in each group, the sum of gradient difference can be bounded by the group size multiplied by the maximum difference in the group, i.e., $|G_i| \max_{k \in G_i} \|\nabla l_k^+(\theta) - \nabla l_{\gamma(\phi_\theta(k))}^+(\theta)\|$. Thus, Eq. 3 can be bounded by the result of Eq. 4. Next, we focus on how to minimize the bound.

Recap from Example 3, intuitively, the bound (i.e., the right hand in Eq. 4) will be minimized when $\phi_\theta(k)$ is set to assign each t_k^+ to the closest tuple in the coreset C, w.r.t. the gradient, as follows:

$$
\sum_{i=1}^{g} |G_i| \max_{k \in G_i} \min_{c_j \in C} \|\nabla l_k^+(\theta) - \nabla l_{\gamma(j)}^+(\theta)\|
\tag{5}
$$

However, given a coreset, it is infeasible to compute Eq. 5 because we have to iterate each tuple in every group, which is equivalent to iterate the entire T^+, but the large T^+ will not be materialized due to the inefficiency. To address this issue, considering Eq. 4 and Eq. 5 and leveraging max-min inequality [7] over Eq. 5, we can get:

$$
\begin{aligned}
&\|\sum_{i=1}^{N} \nabla l_i^+(\theta) - \sum_{j=1}^{|C|} w_j \nabla l_{\gamma(j)}^+(\theta)\| \\
&\leq \sum_{i=1}^{g} |G_i| \max_{k \in G_i} \min_{c_j \in C} \|\nabla l_k^+(\theta) - \nabla l_{\gamma(j)}^+(\theta)\| \\
&\leq \sum_{i=1}^{g} |G_i| \min_{c_j \in C} \max_{k \in G_i} \|\nabla l_k^+(\theta) - \nabla l_{\gamma(j)}^+(\theta)\|
\end{aligned}
\tag{6}
$$

Given Eq. 6, we can iterate the much smaller coreset C and the gradient approximate error can be bounded using the largest gradient difference between $c_j \in C$ and tuples within each group, i.e., $\max_{k \in G_i} \|\nabla l_k^+(\theta) - \nabla l_{\gamma(j)}^+(\theta)\|$, which can be computed efficiently without joining all tables (see § 5). In this way, all tuples within a group will be assigned to the same tuple in the coreset, i.e., $\forall k \in G_i, \phi_\theta(k) = j$.

So far, the deductions we have discussed only consider the case for a particular θ. Obviously, it is prohibitively expensive to explore every possible θ. Next, we illustrate how to bound the gradient approximation error for the parameter space ϑ.

4.2 Upper Bound for the Parameter Space ϑ and Groups

Fortunately, it has been proved in recent works [5, 22, 39] that for convex ML algorithms, e.g., linear regression, logistic regression, the normed gradient difference between tuples can be efficiently bounded by:

$$
\forall i, j, \max_{\theta \in \vartheta} \|\nabla l_i(\theta) - \nabla l_j(\theta)\| \leq \max_{\theta \in \vartheta} O(\|\theta\|) \cdot \|\mathbf{x}_i - \mathbf{x}_j\|
\tag{7}
$$

where $\|\mathbf{x}_i - \mathbf{x}_j\|$ denotes the Euclidean distance between the feature vectors of two tuples. Since $O(\|\theta\|)$ is a constant, we can conclude that **the gradient approximation error can be bounded independent of the optimization problem in practice, i.e., any particular θ**. Thus, based on the results in Eq. 6, we can get:

$$
\begin{aligned}
&\max_{\theta \in \vartheta} \sum_{i=1}^{g} |G_i| \min_{c_j \in C} \max_{k \in G_i} \|\nabla l_k^+(\theta) - \nabla l_{\gamma(j)}^+(\theta)\| \\
&\leq \sum_{i=1}^{g} |G_i| \min_{c_j \in C} \max_{k \in G_i} \max_{\theta \in \vartheta} \|\nabla l_k^+(\theta) - \nabla l_{\gamma(j)}^+(\theta)\| \\
&\leq \underbrace{c}_{const} \cdot \sum_{i=1}^{g} |G_i| \min_{c_j \in C} \max_{k \in G_i} \|\mathbf{x}_k^+ - \mathbf{x}_{\gamma(j)}^+\|
\end{aligned}
\tag{8}
$$

Feature similarity. For ease of representation, we use similarity $s_{ji} = 1 - normalized(\max_{k \in G_i} \|\mathbf{x}_k^+ - \mathbf{x}_{\gamma(j)}^+\|)$ to denote the minimum similarity of feature vectors between $c_j \in C$ and tuples in group $\mathcal{G}_i$.

Obviously, the minimization of $\min_{c_j \in C} \max_{k \in G_i} \|\mathbf{x}_k^+ - \mathbf{x}_{\gamma(j)}^+\|$ is equivalent to the maximization of $\max_{c_j \in C} s_{ji}$.

Therefore, the gradient approximation error can be bounded by $\sum_{i=1}^{g} |G_i| \max_{c_j \in C} s_{ji}$.

Note that Eq. 7 holds for tuples with the same label [5, 22]. Hence, we need to select subsets of coresets for tuples with different labels and combine them. For example, given a binary classification task (30% of label 0 and 70% with label 1), to select a coreset with size K, we separately select a coreset of size 30%K for tuples with label 0 and another coreset of size 70%K for tuples with label 1, and then merge them. For regression tasks, we will cluster tuples with similar labels, select subsets of coresets for these clusters and merge them.

Problem (GGEM). The problem shown in Eq. 2 can be converted to the **group-based gradient approximation error minimization (GGEM)** problem as follows:

$$
C^* = \arg\max_{C \subseteq T^+} \sum_{i=1}^{g} |G_i| \max_{c_j \in C} s_{ji}, \text{ s.t. } |C| \leq K
\tag{9}
$$

At a high level, when $T^+ = T \bowtie T_1 \bowtie \cdots \bowtie T_m$ is too large to compute the coreset, the GGEM problem is to efficiently select an optimal coreset C^* such that the tuples in C^* and their associated weights can well approximate the full gradient in T^+, by grouping T^+ and distributing the computation to multiple tables based on the groups. In § 5, we will show how to compute the partial feature similarity s_{ji} efficiently without the fully materialized T^+. Next, we will prove that even if s_{ji} is known, GGEM problem is NP-hard. We will also discuss its submodular property, based on which a greedy algorithm will be designed in § 5.

THEOREM 1. *The* GGEM *problem is NP-hard.*

PROOF. We start the proof by considering a special case of GGEM problem where $\forall i \in [1, g], |G_i| = 1$. Thus, the number of groups g equals to the number of tuples N. Therefore, the problem becomes $C^* = \arg\max_{C \subseteq T^+} \sum_{i=1}^{N} \max_{c_j \in C} s_{ji} = \arg\min_{C \subseteq T^+} \sum_{i=1}^{N} \min_{c_j \in C} \|\mathbf{x}_i^+ - \mathbf{x}_{\gamma(j)}^+\|, |C| \le K$. Naturally, the K-medoid problem [19] can be reduced to the the special case and thus GGEM is NP-hard. □

THEOREM 2. The GGEM problem has the submodular property.

Due to the space limitation, proof of submodular property is left to the technical report [2].

Our scope. Note that we focus on the convex problems trained with gradient descent because for such problems, the gradient difference can be bounded by the difference between feature vectors. In this situation, regardless of any convex ML algorithm or parameter θ, RECON can select the coreset without training in advance.

4.3 Connection between Groups of the Single Augmented Table and Multiple Tables

As mentioned in § 3.3, the joined result T^+ is generally large in scale, which makes it inefficient to directly select coreset over T^+. To address this, our key idea is to divide the large-scale T^+ into some disjoint groups (like the Groupby), each of which contains tuples that have the same attribute values over one or more predefined attributes. In this way, the gradient computation over tuples in T^+ can be pushed down as a pre-computation step in the corresponding groups of each single table respectively, and further bounded by aggregating the results from multiple tables efficiently.

Remark. In this paper, we consider the join type that introduces redundancy in the data, including both tuple redundancy and feature redundancy [14, 29] (*e.g.,* PK-FK, one (many)-to-many joins). For example, considering $R = S \bowtie T$, for table S, the tuple ratio is denoted by $\frac{n_R}{n_S}$ and feature ratio is $\frac{d_R}{d_S}$, where n represents the number of tuples and d represents the number of features. We assume that the tuple ratio and the feature ratio are larger than 1.

5 RECON ALGORITHM

Theory 2 in § 4, *i.e.,* the GGEM problem has the submodular property, tells us that we can use a greedy algorithm with the approximate ratio $(1 - \frac{1}{e})$ to solve the problem of coreset selection without joining multiple tables. In what follows, we will first overview the algorithm in § 5.1, followed by discussing two important components of the algorithm in § 5.2 and § 5.3 respectively.

5.1 Algorithm Overview

Algorithm. Algorithm 2 takes the base table T, a set $\mathcal{T}$ of tables to be augmented and the coreset size K (can be specified by the user) as input, greedily adds the tuple that brings about the largest utility improvement into the coreset such that a near-optimal coreset is finally output. Note that T^+ denotes the result of joining all above tables together, which will not be materialized.

It first pre-computes the feature vector difference between every two tuples in each table, so as to bound the gradients (line 1). Next, it iteratively adds one tuple to the coreset at a time (the first loop in lines 3-12). At each iteration, the tuple with the largest utility

---

Algorithm 2: RECON Algorithm

Input: The train data T, $\mathcal{T} = \{T_1, T_2, ..., T_m\}$, coreset size K.
Output: A coreset $C \subseteq T^+$, weight $W = \{w_j\}, |C| = |W| = K$.

1 Pre-compute table-wise partial feature similarity difference
$D^+ = \{d_{ij}^h | \forall T_h \in \mathcal{T} \cup \{T\}, \forall t_i^h, t_j^h \in T_h, d_{ij}^h = \|\mathbf{x}_i^h - \mathbf{x}_j^h\|\}$.

2 $C = \emptyset$;

3 **while** $|C| < K$ **do**

4   Sample a subset $\mathcal{S}$ from T^+ using T and $\mathcal{T}$;

5   **for** *each tuple* $t_j \in \mathcal{S}$ **do**

6          $U(C \cup \{t_j\}) = 0$;

7   **for** *each group* $\mathcal{G}_i \in \mathcal{G}$ **do**

8   Compute s_{ji} by aggregating $d_{i'j'}^h \in D^+$ from different tables;

9              $U(C \cup \{t_j\}) \mathrel{+}= |G(i)| \max_{c_j \in C \cup \{t_j\}} s_{ji}$;

10          $U(t_j|C) = U(C \cup \{t_j\}) - U(C)$;

11      $t^* = \arg\max_{t_j \in \mathcal{S}} U(t_j|C)$;

12   Add t^* to C ;

13 **for** $j = 1$ *to* $|C|$ **do**

14      $w_j = \sum_{i=1}^{g} |G(i)| \mathbb{I}[j = \arg\max_{c_{j'} \in C} s_{j'i}]$

15 **return** C, W;

---

among T^+ will be selected (line 11). However, it is rather time-consuming to iterate T^+, and thus we can sample a set $\mathcal{S}$ of joined tuples and select one from them (the second for loop in lines 5-10). Afterwards, to compute the utility of each tuple $t_j \in \mathcal{S}$, the third loop (lines 7-9) iterates each group $\mathcal{G}_i$, computes the feature similarity s_{ji} (line 8), reconsiders whether tuples in $\mathcal{G}_i$ should be assigned to t_j, and updates the utility value (line 9). In short, we can derive the coreset C by repeatedly sampling tuples, computing their utilities, and select the best one as a member of the coreset. Finally, we also have to assign each $c_j \in C$ a weight w_j.

Algorithm 2 consists of the following four key components.

[Component ❶: *Partial feature similarity* pre-computation.] As shown in line 1, we first compute the similarity of partial feature vectors in each individual table, in order to bound the gradient based on Eq. 7. This step is easy to implement but rather important. The motivation and the details will be introduced in § 5.2.

[Component ❷: *Joined tuples sampling.*] We use the sampling method proposed in [58] to sample a set $\mathcal{S} \subset T^+$ of tuples as candidates (line 4). Although T^+ will not be materialized, [58] guarantees that these sampled tuples are uniformly sampled from T^+. Thus, there still has an approximate ratio $1 - \frac{1}{e} - \epsilon$ for the greedy algorithm, as discussed in § 3.2, where $|\mathcal{S}| = (N/K) \cdot log(1/\epsilon)$.

[Component ❸: *Feature similarity computation.*] During the above process, computing s_{ji} is challenging because we do not have T^+. To address this, the high level idea is to aggregate the pre-computation results in Component ❶ along with the join keys, which will be introduced in § 5.3.

[Component ❹: *Weight computation.*] As shown in line 14, w_j equals to the sum of the groups that are assigned to c_j timing the group size, because the tuples within a group will be assigned a single tuple in the coreset.

5.2 Partial Feature Similarity Pre-computation

As discussed above, to solve the GGEM problem, it is significant to compute the feature similarity s_{ji}, *i.e.,* the minimum similarity

Figure 6: Partial feature similarity pre-computation.

of feature vectors between c_j and tuples in the group $\mathcal{G}_i$, so as to bound the gradient difference. The computation is highly related to the Euclidean distance between two feature vectors, as shown in Eq. 8. Although we can obtain one vector of c_j through sampling, the other one in the group is not available because we do not want to materialize T^+ and iterate it. Fortunately, each feature vector in T^+ can be represented by the concatenation of $(m + 1)$ sub-vectors from these m candidate tables as well as the base table, i.e., $\mathbf{x}_i^+ = [\mathbf{x}_i^+[0], \mathbf{x}_i^+[1], ..., \mathbf{x}_i^+[m]]$, where $\mathbf{x}_i^+[0]$ denotes the base table part of feature vector of t_i^+. Intuitively, to capture the feature difference, we can first capture the partial feature similarity inside each table, and then aggregate from multiple tables to compute s_{ji}.

Hence, we push down the computation to each individual table as a pre-computation step, which accelerates the coreset selection much. To be specific, for each table T_k, the feature vector difference of any tuple pair, i.e., $d_{ij}^h = \|\mathbf{x}_i^h - \mathbf{x}_j^h\|$ is computed.

In Figure 6, suppose that the tuples colored green from multiple tables will form a group $\mathcal{G}_2$. In pre-computation, $\|\mathbf{x}_j^+[1] - \mathbf{x}_2^1\|$, $\|\mathbf{x}_j^+[1] - \mathbf{x}_3^1\|$, $\|\mathbf{x}_j^+[0] - \mathbf{x}_3^0\|$, $\|\mathbf{x}_j^+[2] - \mathbf{x}_1^2\|$ and $\|\mathbf{x}_j^+[2] - \mathbf{x}_2^2\|$ have been computed when we want to compute s_{12}, which accelerates the coreset selection much because we do not need to compute the differences over the large scale T^+. Note that we are not going to compute the feature similarity between every two tuples in T^+. Recap that in § 4.2, the similarity s_{ji} to be computed, i.e., $s_{ji} = 1 - normalized(\max_{k \in G_i} \|\mathbf{x}_k^+ - \mathbf{x}_{\gamma(j)}^+\|)$ is to denote the minimum similarity of feature vectors between $c_j \in C$ and tuples in group $\mathcal{G}_i$. Once s_{ji} can be computed, we can use it to bound the gradient approximation error, so in the next section, we will discuss how to aggregate pre-computed partial results to derive s_{ji}.

5.3 Gradient Aggregation for Feature Similarity (s_{ji}) Computation

Next, we describe, for a given $c_j \in C$, how to efficiently compute s_{ji}, which is equivalent to compute the difference, i.e., $\max_{k \in G_i} \|\mathbf{x}_k^+ - \mathbf{x}_{\gamma(j)}^+\|$ for all the groups $\mathcal{G}_i \in \mathcal{G}$ using a dynamic programming (DP) algorithm. We mainly illustrate the algorithm using a concrete example as shown in Figure 7. Before that, we first introduce some necessary notations.

Recap that in § 5.2, each sampled tuple $\mathbf{x}_{\gamma(j)}^+$ can be represented by concatenating $m+1$ sub-vectors from different tables respectively. Hence, for each subvector $\mathbf{x}_{\gamma(j)}^+[h], h \in [0, m]$, we use $\gamma_h(j)$ to denote the index of $\mathbf{x}_{\gamma(j)}^+[h]$ in table T_h. In this way, for all tuples in each table T_u, $d_{\gamma_h(j),u}^h = \|\mathbf{x}_{\gamma_h(j)}^h - \mathbf{x}_u^h\| = \|\mathbf{x}^+[h] - \mathbf{x}_u^h\|, u \in [1, |T_u|]$. Besides, we model these tables as a tree structure, where the root is the table that we group on. We use J_h to denote the set of children table index of T_h. R_u^h denotes the intermediate result of t_u^h (The u-th

tuple in T_h) joining with all the descendants of T_h. Next, we use two examples to clarify these.

EXAMPLE 4. *[Join tree] Suppose that $\mathcal{A} = \{A_0\}$, we have a tree in Figure 7, where T_0 is the root, and T_1, T_3 are leaves. Thus, $J_0 = \{1, 2\}$, $J_2 = \{3\}$ and $J_1 = J_3 = \emptyset$. Besides, T_0 and T_1 can be joined on attribute A_2, where values with the same color can be joined together. For example, the first two tuples in T_0 can be joined with the first tuple in T_1. Based on that, R_1^2 denotes the result (i.e., two tuples) of t_1^2 joining with T_3 (the descendant of T_2). R_1^2 denotes the result (i.e., three tuples) of t_1^1 joining with T_1, T_2 and T_3 (the descendants of T_0).*

[Groups] Suppose that we want to group on attribute A_0 (i.e., the group key). Clearly, both T_0 and T^+ will have 3 groups $\mathcal{G}_1, \mathcal{G}_2$ and $\mathcal{G}_3$. For example, the first group $\mathcal{G}_1$ in T_0 has two tuples, but after joining, there will be six tuples in the group.

Our goal. Given the join tree, group key and pre-computation results (§ 5.2), the goal of our DP algorithm is to compute the feature similarity s_{ji} between each group $\mathcal{G}_i$ and the tuple c_j in the coreset, by aggregating the results from multiple tables along with the join keys, without materializing T^+.

Key observation. The feature similarity computation has the optimal substructure. Hence, at a high level, we can group within each individual table on the attributes to be joined, and then use a dynamic programming algorithm to compute the s_{ji} across the join relations. Taking the $\mathcal{G}_1$ as an example, that is, we want to get s_{11} (the similarity/difference between c_1 and $\mathcal{G}_1$). For the two tuples t_1^0 (t_2^0), we have already known $d_{\gamma_0(1),1}^0$ ($d_{\gamma_0(1),2}^0$). Hence, if we can compute the maximum difference of feature vectors of tuples that can join with t_1^0 (t_2^0) from other tables, which serve as the *optimal substructure*, we can add $d_{\gamma_0(1),1}^0$ ($d_{\gamma_0(1),2}^0$) to the corresponding maximum difference and output two values. Finally s_{11} can be computed by choosing the largest one from the two values.

Specifically, to capture the join relations of tuples, except the base table, for each table T_h, we group the tuples in T_h on the attribute that serves as the key to join with T_h's parent. We use P_v^h to denote the v-th group of tuples in T_h. For example, P_1^2 includes the first two tuples in T_2 and P_2^3 just includes the last tuple in T_3, as shown in Figure 7. Note that for the base table, the groups are constructed based on A_0 rather than the join key because T_0 is the root.

We use $dp[u, h]$ to denote the maximum difference between tuples in R_u^h and c_j's corresponding sub-vectors *w.r.t.* tuples in R_u^h. We then use $DP[v, h]$ to denote the maximum dp value among the v-th group of T_h, i.e., $DP[v, h] = \max_{t_u^h \in P_v^h} dp[u, h]$. Thus, we have:

$$dp[u, h] = d_{\gamma_h(j),u}^h + \sum_{h' \in J(h)} DP[v', h'] \qquad (10)$$

where v' denotes the index of group in $T_{h'}$ that can join with t_u^h. Then we can compute $DP[i, 0], i \in [1, g]$ following Eq. 10. Obviously, we can directly compute s_{ji} based on $DP[i, 0]$.

For ease of discussion, here we just consider equi-join where each tuple can join with at most one group of a table. We will discuss how to support fuzzy join later in this section.

Let us illustrate the algorithm using an example to compute s_{11}.

EXAMPLE 5. *[The DP algorithm] We run the DP algorithm from bottom to up. Initially, we compute the dp values for all tuples in the*
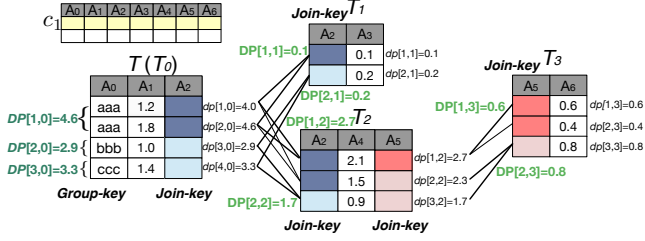
Figure 7: An example of dynamic programming.

leaf nodes, i.e., T_1 *and* T_3 *in Figure 7. For example, to compute* $dp[1, 3]$, *since* $J(3) = \emptyset$, $dp[1, 3] = d^3_{\gamma_3(1),1} = 0.6$ *following Eq. 10. Next we consider how to compute* $DP[1, 3]$. *Group* P^3_1 *consists of* t^3_1 *and* t^3_2 *as they have the same value on the join key attribute* A_5, *so we can compute* $DP[1, 3] = \max(dp[1, 3], dp[2, 3]) = 0.6$. *Thus, the largest in the group (*i.e., $dp[1, 3]$*) will be propagated (marked as the bold line) to its father relation (*i.e., T_2*) for further computation. Then we come to* T_2. *Similarly, to compute* $dp[1, 2]$, *considering Eq. 10* $d^2_{\gamma_2(1),1}$ *is added first. After that, since* $J(2) = 3$ *and* t^2_1 *joins with tuples in* P^3_1, $DP[1, 3]$ *will also be added to* $dp[1, 2]$. *Therefore, we obtain* $dp[1, 2] = d^2_{\gamma_2,1} + DP[1, 3] = 2.1 + 0.6 = 2.7$. *After that,* $DP[1, 2]$ *is computed as* $\max(dp[1, 2], dp[2, 2]) = 2.7$, *which is then propagated to the first two tuples in* T_0 *that join with tuples in* P^2_1. *Finally, we come to* T_0. *To compute* $dp[1, 0]$, $d^0_{\gamma_0(1),1}$ *is first added to* $dp[1, 0]$ *as well. Since* $J(0) = \{1, 2\}$, $DP[1, 1]$ *and* $DP[1, 2]$ *will be added to* $dp[1, 0]$. *Therefore, we can obtain* $dp[1, 0] = d^0_{\gamma_0(1),1} + DP[1, 1] + DP[1, 2] = 1.2 + 0.1 + 2.7 = 4.0$. *Finally,* $DP[1, 0]$ *is computed as* $DP[1, 0] = \max(dp[1, 0], dp[2, 0]) = 4.6$.

Complexity Analysis of Algorithm 2. Recap that $|T^+| = N$. For ease of illustration, we use $O(n)$ to denote the average size of each single table, *i.e.*, $O(|T_h|) = O(n)$. Besides, we use D (d) to denote the number of feature dimensions of T^+ (T_h on average) respectively. Next, we analyze the time complexity from two aspects, *i.e.*, the pre-computation and greedy algorithm with 3 loops.

Partial Feature Similarity Pre-computation. In this stage, we need to pre-compute the difference of feature vectors between every two tuples in each table, so the time complexity is $O(n^2d)$ when the number of tables can be regarded as a constant.

Greedy Coreset Selection. To select the coreset C, the greedy algorithm repeats K times. Each time a set of tuples $\mathcal{S} \subset T^+$ is sampled. To get i.i.d. uniform samples, we use the exact weight algorithm from Zhao et al. [58], which has no accept-reject step in the sampling phase. After a pre-computation in $O(n)$, we can get a sample from T^+ in $O(1)$ [51, 53, 58]. For each element $t_j \in \mathcal{S}$, we need to compute $U(t_j|C)$, which involves the computation of s_{ji} for every group $\mathcal{G}_i \in \mathcal{G}$ using the DP algorithm. The complexity of DP is linear to the total sizes of relations, *i.e.*, $O(|T_0|+|T_1|+\cdots+|T_m|) = O(n)$. For ease of representation, we use S to denote $|\mathcal{S}|$. Therefore, The time complexity of this part is $O(K \cdot |\mathcal{S}| \cdot n) = O(nKS)$.

Total Time Complexity. In summary, the total time complexity of our approach is $O(n^2d + nKS)$.

To show the superiority of our method, we also illustrate the time complexity of the SOTA single-table coreset selection algorithm

on T^+, if the join result is materialized. First, they compute feature vector differences between every two tuples in T^+, so as to bound the gradient, leading to a time complexity of $O(N^2D)$. Afterwards, to compute the utility of each tuple , their methods have to iterate every tuple in T^+, leading to a time complexity of $O(K \cdot S \cdot N) = O(NKS)$. In total, the time complexity is $O(N^2D + NKS)$.

For feature-enrich ML, $N \gg n$ always holds when various types of joins exist. Thus, our method can much improve the efficiency.

Discussion. For fuzzy join, each tuple t^h_u may join with multiple groups $P^{h'}_{v'}$ for $h' \in J(h)$ of a table. We can extend RECON to handle this by changing the $DP[v', h']$ in Eq. 10 into $dp[u, h] = d^h_{\gamma_h(j),u} + \sum_{h' \in J(h)} \max_{v' \in V} DP[v', h']$, where V represents all the groups in $T_{h'}$ that can join with t^h_u. For the grouping key, if the attributes in $\mathcal{A}$ specified by the user are distributed in multiple tables, we can run our algorithm by randomly selecting a table from these tables as the root and aggregating the results using the DP algorithm.

Note that the upper bound of the gradient difference derived in § 4 only holds for points with similar labels. Thus, theoretically we need to select subsets separately. The above analysis assumes that all tuples correspond to same labels. However, in practice, data has different labels. Therefore, data with a certain label generally represents only a small fraction of the total data. The total amount of calculation is much less than the above complexity.

Convergence Rate Analysis of Algorithm 2. In the field of ML, convergence rate reflects how fast the machine learning algorithm can find the optimal parameters. The higher the convergence rate, the fewer iterations the model needs to converge. Specifically, we can compute the convergence rate by comparing the parameter θ computed in the k-th and the $(k + 1)$-th iteration to the optimal one. The detailed proof is left to the technical report [2].

As proved in [2], the convergence rate of Algorithm 2 is at the same rate of $O(\frac{1}{\sqrt{k}})$ as the convergence rate for incremental gradient descent on the full data T^+ [42]. Hence, theoretically, RECON needs the same number of epochs to converge as training on the full data. In this situation, since the coreset is much smaller than the full data, the efficiency is much improved.

6 EXPERIMENTS

The key questions we seek to answer are: (1) How does RECON perform to select a well-performed coreset with an appropriate size, as an end-to-end solution (§ 6.2)? (2) What about the effectiveness and efficiency of RECON, compared with baselines (§ 6.3 - § 6.5)?

6.1 Experimental Settings

Dataset. We used 5 widely-used real-world datasets that covered various data characteristics, *e.g.*, the dataset size varying from the magnitude of 10^4 to 10^7. Table 1 shows the statistics of the datasets.

(1) Brazil [1] is a dataset with a multi-classification task to predict *"the review score of an order given by the customer"* with four tables.

(2) IMDB [32] is a dataset that *"predicts the score of movies"* with 7 tables. Obviously, we can regard it as a regression task to predict the score. To show more thorough experiments, similar to [14], we also regard it as a classification task by dividing the rating scores in to 5 equal intervals (*i.e.*, grades) and predict the grade.

Table 1: Statistics of datasets.

Dataset	# Tables	# Rows (T^+)	# features (T^+)	Task
Brazil	4	98,463	9	Class.
IMDB	7	674,466	41	Class./Reg.
IMDB-Large	7	21,303,410	41	Class./Reg.
Stack	3	6,347,553	178	Reg.
Taxi	5	2,792,376	30	Reg.

(3) IMDB-Large is similar to IMDB, except that IMDB-Large uses all tuples in Cast_info, producing 21,303,410 tuples for T^+.

(4) Stack [37] contains questions and answers from the StackExchange, which *"predicts the reputation of users"* as a regression task.

(5) Taxi [15] is to *"predict the number of vehicle collisions in New York City for each day"*. In particular, we have *fuzzy join* on this dataset, *e.g.,* the Weather table can join with the base table on the attribute *w.r.t.* time, but the weather data is represented by the granularity of minutes, hours, or days.

Following [30], for every dataset, the base table is randomly shuffled and divided into 50%/25%/25% proportions as train/validation/test set. All other tables will be used for feature augmentation via joins while training, validating and testing.

The group key $\mathcal{A}$ of each dataset is specified by the user. Although the user can specify any set of attributes as the group key, by default, we use the (primary) key of the base table, if available. We used {review_id}, {movie_id}, {movie_id}, {user_id} and {datetime} as group keys for Brazil, IMDB, IMDB-Large, Stack and Taxi in their base tables respectively.

Baselines. We compared with several baselines.
(1) Base uses the base table T as train data to train ML models (see *e.g.,* Figure 3(a)–❶).
(2) Full uses the fully augmented table T^+ as the train data to train ML models (see *e.g.,* Figure 3(a)–❸).
(3) Sample-Join [58] uniformly samples tuples from T^+ as train data without materializing the join result.
(4) Join-Coreset [39, 41] selects the coreset over fully materialized join result T^+ and uses the coreset as train data (see *e.g.,* Figure 3(a)–❺). We use the popular single-table coreset selection algorithm [39] that follows the paradigm in Figure 5(b).
(5) Coreset-Join first selects a coreset from the base table T, then joins with tables in $\mathcal{T}$ and finally trains on the join result.
(6) FML [14, 29, 49] (factorized ML) focuses on accelerating batch gradient descent algorithm by decomposing the ML computations through joins. Among these methods, [14] is a general one for different ML algorithms, so we compare with it in § 6.5. In other sections, we focus on the stochastic gradient descent algorithm that is widely used in practice due to its high efficiency.

Hyper-parameter Setting. For the classification task and regression task, we train logistic regression and linear regression models by default respectively, where L2-regularization (regularization coefficient=10^{-5}) and stochastic gradient descent (SGD) are applied. The influence of using different ML models will be evaluated in § 6.3. For training, we fix the number of training epochs to 20. We use *k-inverse decay scheduling, i.e.,* $\alpha_k = \alpha_0/(1+bk)$, where α_0 and b are tuned as hyperparameters independently for different methods. The sample size S is set to 500.

Evaluation Metrics. We evaluate the efficiency in an end-to-end way, including both the consuming time of coreset selection and
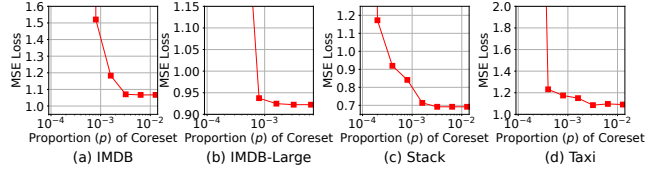

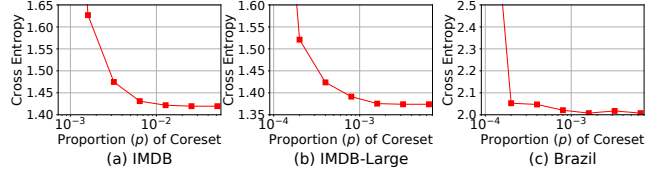**Figure 8: End-to-end coreset selection for regression.**

Figure 9: End-to-end coreset selection for classification.

model training. For effectiveness, we use different evaluation metrics for different tasks. For classification tasks, following previous works [15, 58], we use *model prediction accuracy* as evaluation metric. For regression tasks, following [29], we use *root mean squared error* $(RMSE = \sqrt{\frac{\sum_{t=1}^{N}(\hat{y}_t - y_t)^2}{N}})$ as evaluation metric.

6.2 End-to-end Coreset Selection

Recap that in § 5, the proposed algorithm takes as input a user-specified K as the size of the coreset. To realize an end-to-end solution, one may consider how to choose an appropriate size of the coreset. In this part, we propose a simple yet effective approach to achieve this. For ease of explanation, we introduce $p = \frac{K}{N}$ to denote the proportion of coreset compared to the full data in size. N, *i.e.,* $|T^+|$ can be computed efficiently using [58] before it is materialized.

Approach. We start from a coreset in a small size, train on it and evaluate on the validation set, enlarge the coreset and iteratively train until the performance cannot improve much. Specifically, we start from $K = 10^{-4}N$, *i.e.,* $p = 10^{-4}$ and train an initial model. Then, we iteratively enlarge the coreset by 2 times and train. To evaluate each coreset, we apply the model on the validation set and compute the validation loss. If the loss decreases and remains stable within several successive iterations, we stop enlarging the coreset.

Validation loss. Figure 8 -9 show the validation loss (*i.e.,* MSE loss for regression task and cross entropy loss for classification task on the y-axis) by varying the coreset size (the x-axis). At a high level, with the number of tuples of a coreset increasing, the validation loss decreases rapidly first and then remains stable. As Figure 8(c) shows, on Stack dataset, when $K = 0.0032 \times N = 10156$, the loss is 0.69 and then it does not decrease much. We set that within three successive iterations, if the loss varies no more than 1%, we can stop. So finally, we can return the coreset with a size of 10156.

Efficiency. One may consider whether the end-to-end coreset selection including iterative training is time-consuming. The answer is No. More concretely, with the coreset size increasing, it obviously spends more time because both the iterative training and coreset selection consume time, but it is still efficient. For example, it takes 13.3 mins, 4.6 mins, and 1.1 mins on the task of IMDB-Large and Taxi for regression, and Brazil for classification respectively to perform the end-to-end coreset selection. However, if train on T^+

Figure 10: Efficiency. (a,b,c,d): Regression tasks; (e,f,g): Classification tasks.


Figure 11: Effectiveness. (a,b,c,d): Regression tasks; (e,f,g): Classification tasks.


Figure 12: Training loss comparison for IMDB and Stack.

Figure 13: Effectiveness and efficiency by varying p.

Figure 14: Compare with FML.

(*i.e.,*Full), it requires 782 mins, 72 mins, and 18 mins respectively. The reasons are (1) The size of coreset is small and thus efficient to train. (2) The number of iterations is not large, and we can fine-tune the model in last iteration without training from scratch. (3) Coreset selection algorithm is efficient and can be done incrementally.

Summary. This end-to-end coreset selection approach provides a way to select an appropriate coreset size. Although several iterations are needed, it is also efficient because the coreset size is small and the coreset selection process is fast.

6.3 Comparison with Baselines

Using the coreset size of each dataset selected in § 6.2, we compare efficiency and effectiveness with baselines. To achieve a fair comparison, for the baselines that sample a subset of tuples to train, we sample the same number of tuples as the coreset size. For Coreset-Join, we set the proportion of coreset over the base table with the same p as RECON.

Efficiency. We show the total time including both coreset selection and model training for different baselines and RECON in Figure 10, given the coreset with the best size. We can see that in general, RECON achieves efficiency improvement nearly two orders of magnitudes compared with Full and Join-Coreset. For example, on

dataset IMDB-Large for regression, RECON takes 13.3 minutes, which is nearly 2 orders of magnitudes more efficient than Full (782 mins) and Join-Coreset (612 mins). For classification, on dataset IMDB-Large, RECON takes 110 min, which is also almost 2 orders of magnitudes more efficient than Full (2.8 days) and Join-Coreset (2.2 days). In addition, on dataset Stack and Taxi for regression, RECON takes 11.8 mins and 4.6 mins respectively, which is still an order of magnitude faster than Full (4 hours, 72 mins) and more efficient than Join-Coreset (0.5 hours, 9 mins) over 2 times. The reason is that Full has to train over a large amount of training tuples, *i.e.,* T^+. Although Join-Coreset trains over a coreset with the same size of RECON, it selects the coreset based on T^+, which is rather inefficient. RECON outperforms them because it computes the coreset directly from these tables to be joined. Besides, RECON takes a slightly longer time than other baselines, *e.g.,* Sample-Join (40 mins), Coreset-Join (35 mins) and Base (11 mins) on IMDB-Large for classification, because Base does not need to augment features and select the coreset, Sample-Join just uniformly samples without considering the gradients and Coreset-Join computes the coreset over the much smaller base table T. But they cannot achieve high accuracy as discussed next.

Effectiveness. For regression tasks, on dataset IMDB-Large, we find in § 6.2 that $p = 0.0016$ is the best choice. In this situation, we can observe in Figure 11(b) that RECON has an RMSE of 0.998, which outperforms Sample-Join (1.546). The reason is that Sample-Join just samples for training without considering the gradient approximation. RECON outperforms Base because more useful features are augmented. Besides, RECON also outperforms Coreset-Join (1.603)

Table 2: Convergence results for regression

Dataset	# iterations to converge		Test RMSE after convergence	
	RECON	Full	RECON	Full
IMDB	37,000	6,200,700	1.025	1.023
IMDB-Large	330,000	201,300,000	0.998	0.985
Stack	174,700	62,912,000	0.749	0.747
Taxi	72,000	20,149,000	0.761	0.760

Table 3: Convergence results for classification

Dataset	# iterations to converge		Test accuracy after convergence	
	RECON	Full	RECON	Full
IMDB	20,500	6,401,000	0.662	0.664
IMDB-Large	340,000	201,220,000	0.679	0.681
Brazil	780	874,800	0.606	0.607

because the selected coreset of Coreset-Join does not consider features to be augmented. Furthermore, RECON almost has the same RMSE as Join-Coreset (0.988) and Full (0.985) because RECON can well approximate the full gradient accurately with theoretical guarantees. That is, training on the coreset (only 0.0016 proportion of the full data) can achieve the almost same performance as training over the full data. For classification tasks, we have similar observations. For example, on dataset IMDB-Large, we can observe from Figure 11(f) that RECON has an accuracy of 0.679, which is higher than Base (0.607), Coreset-Join (0.628), Sample-Join (0.637), and also close to Join-Coreset (0.679) and Full (0.681).

Loss. We show the training loss for IMDB of classification task and Stack of regression task in Figure 12. We can see that RECON converges to almost the same loss as Full, which demonstrates that RECON can accurately estimate the gradient with theoretical guarantees, and thus achieve the same performance as the full data.

Varying the coreset size. We also evaluate the effectiveness and efficiency by varying p in Figure 13 on IMDB-Large of regression task. We only compare RECON with Sample-Join and Join-Coreset, because only they can generate training data of different given sizes. The result of Full is also plotted as a comparison.

Figure 13(a) shows that RMSE of RECON decreases rapidly first, and then remains stable when approaching the best coreset size. RECON outperforms Sample-Join a lot and almost has the same performance as Join-Coreset on all p because RECON can approximate the full gradient well. For efficiency, in Figure 13(b), RECON is more than one order of magnitude faster than Join-Coreset because the coreset selection of RECON has a lower time complexity than Join-Coreset. RECON is only a little slower than Sample-Join because coreset selection of RECON often takes up a small proportion of time compared with iterative training.

Summary. RECON achieves much acceleration (because it computes the coreset without fully materializing the augmented table) for feature-rich ML without sacrificing much effectiveness (because it has the theoretical guarantee on the gradient computation). In addition, we also test on non-convex models and find similar observations, although there is no theoretical guarantees about the gradient. Due to the space limitation, the results are reported in [2].

6.4 Convergence Evaluation

In § 5.3, we have theoretically proved the convergence rate of RECON. To empirically test the convergence of different methods, we compute test performance with the increase of SGD iterations. The concrete numbers of iterations to converge and the converged test performance are reported in Table 2 and Table 3. For all the datasets, training on coresets (RECON) converges much faster than Full. From the results in Table 2 and Table 3, we can observe

that the speedup of RECON is generally more than two orders of magnitudes. For example, on Stack for regression, training on the coreset (RECON, only 0.0032 proportion of the full data) converges in 174,700 iterations, which is 360 times faster than Full (62,912,000 iterations). In addition, the speedup does not affect the test performance after convergence much, *e.g.,* on IMDB for classification, the test accuracy after convergence of RECON (0.662) is similar to Full (0.664). RECON converges fast with high accuracy, because the coreset selected by RECON is much smaller than the full data, while still approximating the full gradient with theoretical bound.

6.5 Comparison with FML

FML only supports batch gradient training, so we also use batch gradient descent to train for a fair comparison. Since FML aims to accelerate the training process over the full data, it has the same performance as Full, so we only compare the efficiency with FML. We compare the total time including both coreset selection and model training between RECON and FML on Brazil and Stack. Figures 14(a)-(b) report the result, which shows that RECON outperforms FML on both datasets. This is because although FML improves the efficiency by reducing the linear algebra computations, it still needs training over the full data, while our method trains over the judiciously selected coreset with a small size.

6.6 Additional Experiments

Note that joins may change the original distribution of the base table T, *i.e.,* one tuple in T may correspond to multiple tuples in T^+, and thus the effectiveness can be evaluated both on T or T^+. In the paper, we only report the results on the distribution of the joined table T^+, and leave the results on T in the technical report [2]. We also evaluated and discussed the effect of changing ML models (We put it in our technical report [2]).

7 CONCLUSION

We propose RECON for selecting a coreset of train tuples from an augmented table without materializing it through joins. RECON solves the problem that coreset selection over a big augmented table is time-consuming. The coreset can speed up iterative gradient methods for training ML models (*i.e.,* data-efficient). The augmented features can improve the accuracy of trained ML models (*i.e.,* feature-rich). Extensive experiments verified RECON can much improve the efficiency of coreset selection without sacrificing the performance.

REFERENCES

[1] 2022. https://www.kaggle.com/datasets/olistbr/brazilian-ecommerce/. Accessed: 2022-04-28.
[2] 2022. Coresets over Multiple Tables for Feature-rich and Data-efficient Machine Learning [Technical Report]. https://github.com/for0nething/RECON-TR/blob/main/main.pdf. Last accessed: 2022-09-15.
[3] Amina Adadi. 2021. A survey on data-efficient algorithms in big data era. *J. Big Data* 8, 1 (2021), 1–54.
[4] Omar Y. Al-Jarrah, Paul D. Yoo, Sami Muhaidat, George K. Karagiannidis, and Kamal Taha. 2015. Efficient Machine Learning for Big Data: A Review. *Big Data Res.* 2, 3 (2015), 87–93.
[5] Zeyuan Allen-Zhu, Yang Yuan, and Karthik Sridharan. 2016. Exploiting the structure: Stochastic gradient methods using raw clusters. *Advances in Neural Information Processing Systems* 29 (2016).
[6] Matthias Boehm and Michael Dusenberry et.al. 2016. SystemML: Declarative Machine Learning on Spark. *Proc. VLDB Endow.* 9, 13 (2016), 1425–1436. https://doi.org/10.14778/3007263.3007279
[7] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. 2004. *Convex optimization.* Cambridge university press.
[8] Vladimir Braverman, Dan Feldman, and Harry Lang. 2016. New Frameworks for Offline and Streaming Coreset Constructions. *CoRR* abs/1612.00889 (2016).
[9] Trevor Campbell and Tamara Broderick. 2018. Bayesian Coreset Construction via Greedy Iterative Geodesic Ascent. In *ICML 2018*, Vol. 80. PMLR, 697–705.
[10] Chengliang Chai, Lei Cao, Guoliang Li, Jian Li, Yuyu Luo, and Samuel Madden. 2020. Human-in-the-loop Outlier Detection. In *SIGMOD 2020*. ACM, 19–33. https://doi.org/10.1145/3318464.3389772
[11] Chengliang Chai, Guoliang Li, Jian Li, Dong Deng, and Jianhua Feng. 2016. Cost-Effective Crowdsourced Entity Resolution: A Partial-Order Approach. In *SIGMOD 2016*. ACM, 969–984. https://doi.org/10.1145/2882903.2915252
[12] Chengliang Chai, Jiabin Liu, Nan Tang, Guoliang Li, and Yuyu Luo. 2022. Selective Data Acquisition in the Wild for Model Charging. *Proc. VLDB Endow.* 15, 7 (2022), 1466–1478. https://www.vldb.org/pvldb/vol15/p1466-li.pdf
[13] Chengliang Chai, Jiayi Wang, Yuyu Luo, Zeping Niu, and Guoliang Li. 2022. Data management for machine learning: A survey. *IEEE Transactions on Knowledge and Data Engineering* (2022).
[14] Lingjiao Chen, Arun Kumar, Jeffrey Naughton, and Jignesh M Patel. 2017. Towards Linear Algebra over Normalized Data. *Proceedings of the VLDB Endowment* 10, 11 (2017).
[15] Nadiia Chepurko, Ryan Marcus, Emanuel Zgraggen, Raul Castro Fernandez, Tim Kraska, and David R. Karger. 2020. ARDA: Automatic Relational Data Augmentation for Machine Learning. *Proc. VLDB Endow.* 13, 9 (2020), 1373–1387.
[16] Kai Lai Chung. 1954. On a stochastic approximation method. *The Annals of Mathematical Statistics* (1954), 463–483.
[17] Jeffrey Cohen, Brian Dolan, Mark Dunlap, Joseph M. Hellerstein, and Caleb Welton. 2009. MAD Skills: New Analysis Practices for Big Data. *Proc. VLDB Endow.* 2, 2 (2009), 1481–1492. https://doi.org/10.14778/1687553.1687576
[18] Dan Feldman. 2020. Introduction to Core-sets: an Updated Survey. *CoRR* abs/2011.09384 (2020).
[19] Michael R Garey and David S Johnson. 1979. *Computers and intractability.* Vol. 174. freeman San Francisco.
[20] Isabelle Guyon and André Elisseeff. 2003. An Introduction to Variable and Feature Selection. *J. Mach. Learn. Res.* 3 (2003), 1157–1182.
[21] Joseph M. Hellerstein, Christopher Ré, Florian Schoppmann, Daisy Zhe Wang, Eugene Fratkin, Aleksander Gorajek, Kee Siong Ng, Caleb Welton, Xixuan Feng, Kun Li, and Arun Kumar. 2012. The MADlib Analytics Library or MAD Skills, the SQL. *Proc. VLDB Endow.* 5, 12 (2012), 1700–1711.
[22] Thomas Hofmann, Aurelien Lucchi, Simon Lacoste-Julien, and Brian McWilliams. 2015. Variance reduced stochastic gradient descent with neighbors. *Advances in Neural Information Processing Systems* 28 (2015).
[23] Jiawei Huang, Ruomin Huang, Wenjie Liu, Nikolaos M. Freris, and Hu Ding. 2021. A Novel Sequential Coreset Method for Gradient Descent Algorithms. In *ICML 2021*, Vol. 139. PMLR, 4412–4422.
[24] Rishabh K. Iyer and Jeff A. Bilmes. 2013. Submodular Optimization with Submodular Cover and Submodular Knapsack Constraints. In *NeurIPS 2013*. 2436–2444.
[25] David Justo, Shaoqing Yi, Lukas Stadler, Nadia Polikarpova, and Arun Kumar. 2021. Towards a polyglot framework for factorized ML. *Proc. VLDB Endow.* 14, 12 (2021), 2918–2931.
[26] KrishnaTeja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, and Rishabh K. Iyer. 2021. GLISTER: Generalization based Data Subset Selection for Efficient and Robust Learning. In *AAAI 2021,*. AAAI Press, 8110–8118.
[27] Katrin Kirchhoff and Jeff A. Bilmes. 2014. Submodularity for Data Selection in Machine Translation. In *EMNLP 2014*. ACL, 131–141.
[28] Arun Kumar, Mona Jalal, Boqun Yan, Jeffrey F. Naughton, and Jignesh M. Patel. 2015. Demonstration of Santoku: Optimizing Machine Learning over Normalized Data. *Proc. VLDB Endow.* 8, 12 (2015), 1864–1867.
[29] Arun Kumar, Jeffrey Naughton, and Jignesh M Patel. 2015. Learning generalized linear models over normalized data. In *SIGMOD 2015*. 1969–1984.

[30] Arun Kumar, Jeffrey F. Naughton, Jignesh M. Patel, and Xiaojin Zhu. 2016. To Join or Not to Join?: Thinking Twice about Joins before Feature Selection. In *SIGMOD 2016*. ACM, 19–34.
[31] Hai Lan, Zhifeng Bao, and Yuwei Peng. 2021. A Survey on Advancing the DBMS Query Optimizer: Cardinality Estimation, Cost Model, and Plan Enumeration. *Data Sci. Eng.* 6, 1 (2021), 86–101.
[32] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter A. Boncz, Alfons Kemper, and Thomas Neumann. 2015. How Good Are Query Optimizers, Really? *Proc. VLDB Endow.* 9, 3 (2015), 204–215.
[33] Guoliang Li, Chengliang Chai, Ju Fan, Xueping Weng, Jian Li, Yudian Zheng, Yuanbing Li, Xiang Yu, Xiaohang Zhang, and Haitao Yuan. 2017. CDB: Optimizing Queries with Crowd-Based Selections and Joins. In *SIGMOD 2017*. ACM, 1463–1478. https://doi.org/10.1145/3035918.3064036
[34] Side Li, Lingjiao Chen, and Arun Kumar. 2019. Enabling and Optimizing Non-Linear Feature Interactions in Factorized Linear Algebra. In *SIGMOD 2019*. 1571–1588.
[35] Jiabin Liu, Chengliang Chai, Yuyu Luo, Yin Lou, Jianhua Feng, and Nan Tang. 2022. Feature Augmentation with Reinforcement Learning. In *ICDE 2022, Kuala Lumpur, Malaysia, May 9-12, 2022*. IEEE, 3360–3372.
[36] Jiabin Liu, Fu Zhu, Chengliang Chai, Yuyu Luo, and Nan Tang. 2021. Automatic Data Acquisition for Deep Learning. *Proc. VLDB Endow.* 14, 12 (2021), 2739–2742. https://doi.org/10.14778/3476311.3476333
[37] Ryan Marcus, Parimarjan Negi, Hongzi Mao, Nesime Tatbul, Mohammad Alizadeh, and Tim Kraska. 2021. Bao: Making Learned Query Optimization Practical. In *SIGMOD 2021*. 1275–1288.
[38] Baharan Mirzasoleiman, Ashwinkumar Badanidiyuru, Amin Karbasi, Jan Vondrák, and Andreas Krause. 2015. Lazier than lazy greedy. In *AAAI*, Vol. 29.
[39] Baharan Mirzasoleiman, Jeff A. Bilmes, and Jure Leskovec. 2020. Coresets for Data-efficient Training of Machine Learning Models. In *ICML 2020*, Vol. 119. 6950–6960.
[40] Baharan Mirzasoleiman, Kaidi Cao, and Jure Leskovec. 2020. Coresets for Robust Training of Deep Neural Networks against Noisy Labels. In *NeurIPS 2020*.
[41] Alexander Munteanu and Chris Schwiegelshohn. 2018. Coresets-methods and history: A theoreticians design pattern for approximation and streaming algorithms. *KI-Künstliche Intelligenz* 32, 1 (2018), 37–53.
[42] Angelia Nedić and Dimitri Bertsekas. 2001. Convergence rate of incremental subgradient algorithms. In *Stochastic optimization: algorithms and applications*. Springer, 223–264.
[43] Dan Olteanu and Maximilian Schleich. 2016. F: Regression Models over Factorized Views. *Proc. VLDB Endow.* 9, 13 (2016), 1573–1576.
[44] Xuedi Qin, Chengliang Chai, Yuyu Luo, Nan Tang, and Guoliang Li. 2020. Interactively Discovering and Ranking Desired Tuples without Writing SQL Queries. In *SIGMOD2020*. ACM, 2745–2748.
[45] Xuedi Qin, Chengliang Chai, Yuyu Luo, Tianyu Zhao, Nan Tang, Guoliang Li, Jianhua Feng, Xiang Yu, and Mourad Ouzzani. 2021. Ranking Desired Tuples by Database Exploration. In *ICDE 2021*. IEEE, 1973–1978.
[46] LKPJ Rdusseeun and P Kaufman. 1987. Clustering by means of medoids. In *Proceedings of the statistical data analysis based on the L1 norm conference, neuchatel, switzerland*, Vol. 31.
[47] Steffen Rendle. 2013. Scaling Factorization Machines to Relational Data. *Proc. VLDB Endow.* 6, 5 (2013), 337–348.
[48] Yuji Roh, Kangwook Lee, Steven Euijong Whang, and Changho Suh. 2021. FairBatch: Batch Selection for Model Fairness. In *ICLR 2021*. OpenReview.net.
[49] Maximilian Schleich, Dan Olteanu, and Radu Ciucanu. 2016. Learning Linear Regression Models over Factorized Joins. In *SIGMOD 2016*. ACM, 3–18.
[50] Vraj Shah, Arun Kumar, and Xiaojin Zhu. 2017. Are Key-Foreign Key Joins Safe to Avoid when Learning High-Capacity Classifiers? *Proc. VLDB Endow.* 11, 3 (2017), 366–379.
[51] Jiayi Wang, Chengliang Chai, Jiabin Liu, and Guoliang Li. 2021. FACE: A Normalizing Flow based Cardinality Estimator. *Proc. VLDB Endow.* 15, 1 (2021), 72–84.
[52] Ruofan Wu, Feng Zhang, Jiawei Guan, Zhen Zheng, Xiaoyong Du, and Xipeng Shen. 2022. DREW: Efficient Winograd CNN Inference with Deep Reuse. In *WWW (WWW '22)*. Association for Computing Machinery, 1807–1816.
[53] Zongheng Yang, Amog Kamsetty, Sifei Luan, Eric Liang, Yan Duan, Xi Chen, and Ion Stoica. 2020. NeuroCard: One Cardinality Estimator for All Tables. *Proc. VLDB Endow.* 14, 1 (2020), 61–73. https://doi.org/10.14778/3421424.3421432
[54] Haitao Yuan and Guoliang Li. 2021. A Survey of Traffic Prediction: from Spatio-Temporal Data to Intelligent Transportation. *Data Sci. Eng.* 6, 1 (2021), 63–85.
[55] Feng Zhang, Jidong Zhai, Xipeng Shen, Onur Mutlu, and Xiaoyong Du. 2022. POCLib: A high-performance framework for enabling near orthogonal processing on compression. *TPDS* 33, 2 (2022), 459–475.
[56] Bo Zhao and Hakan Bilen. 2021. Dataset condensation with differentiable siamese augmentation. In *ICML*. PMLR, 12674–12685.
[57] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. 2021. Dataset Condensation with Gradient Matching. *ICLR* 1, 2 (2021), 3.
[58] Zhuoyue Zhao, Robert Christensen, Feifei Li, Xiao Hu, and Ke Yi. 2018. Random Sampling over Joins Revisited. In *SIGMOD 2018*. ACM, 1525–1539.