

Aggregate-based Training Phase for ML-based Cardinality Estimation

Lucas Woltmann¹, Claudio Hartmann¹, Dirk Habich¹, Wolfgang Lehner¹



Abstract: Cardinality estimation is a fundamental task in database query processing and optimization. As shown in recent papers, machine learning (ML)-based approaches may deliver more accurate cardinality estimations than traditional approaches. However, a lot of training queries have to be executed during the *model training phase* to learn a data-dependent ML model making it very time-consuming. Many of those training or example queries use the same base data, have the same query structure, and only differ in their selective predicates. To speed up the model training phase, our core idea is to determine a *predicate-independent pre-aggregation* of the base data and to execute the example queries over this pre-aggregated data. Based on this idea, we present a specific *aggregate-based training phase* for ML-based cardinality estimation approaches in this paper. As we are going to show with different workloads in our evaluation, we are able to achieve an average speedup of 63 with our *aggregate-based training phase* and thus outperform indexes.

Keywords: cardinality estimation; machine learning; database support; pre-aggregation

1 Introduction

Due to skew and correlation in data managed by database systems (DBMS), query optimization is still an important challenge. The main task of query optimization is to determine an efficient execution plan for every SQL query, whereby most of the optimization techniques are cost-based [Le15]. For these techniques, cardinality estimation has a prominent position with the task to approximate the number of returned tuples for every query operator within a query execution plan [HN17, Le15, MNS09, YW79]. Based on these estimations, various decisions are made by different optimization techniques such as choosing (i) the right join order [FM11], (ii) the right physical operator variant [Ro15], (iii) the best-fitting compression scheme [Da19], or (iv) the optimal operator placement within heterogeneous hardware [KHL17]. However, to make good decisions in all cases, it is important to have cardinality estimations with high accuracy.

As shown in recent papers [Ki19b, Li15], including our own work [Wo19b], machine learning-based cardinality estimation approaches are able to meet higher accuracy requirements, especially for highly correlated data. While traditional approaches such as histogram-based and frequent values methods assume data independence for their estimation [Le15], ML-based approaches assume that a sufficiently deep neural network can

¹ Technische Universität Dresden, Database Systems Group, 01062 Dresden, Germany,
firstname.lastname@tu-dresden.de

model the very complex data dependencies and correlations [Ki19b]. For this reason, ML-based cardinality estimation approaches may thus give much more accurate estimations as clearly demonstrated in [Ki19b, Li15, Wo19b]. However, the main drawback of these ML-based techniques compared to traditional approaches is the high construction cost of the *data-dependent* ML-models based on the underlying supervised learning approach. During the so-called *training phase*, the task of *supervised learning* is to train a model, or more specifically learn a function, that maps input to an output based on example (input, output) pairs. Thus, in the case of cardinality estimation, many pairs consisting of (query, output-cardinality) are required during the *training phase*. To determine the correct output-cardinalities, the queries have to be executed [Ki19b, Wo19b], whereby the execution of those example queries can be very time consuming, especially for databases with many tables, many columns, and millions or billions of tuples resulting in a heavy load on the database system.

Core Contribution. To overcome these shortcomings, we propose a novel training phase *based on pre-aggregated data* for ML-based cardinality estimation approaches. Usually, as described in [Ki19b, Wo19b], every example query is (i) rewritten with a count aggregate to retrieve the correct output-cardinality and (ii) executed individually. However, many of those example queries use the same base data, have the same query structure, and only differ in their selective predicates. To optimize the query execution, our core idea is to provide a *predicate-independent pre-aggregation* of the base data and to execute the example queries over this pre-aggregated data. Consequently, the set of similar example queries has to read and process less data because the *pre-aggregation* is a compact representation of the base data. To realize this *pre-aggregation*, the most common solution in DBMS is to create a *data cube* for storing and computing aggregate information [Gr96]. However, this *pre-aggregation* is only beneficial if the execution of the example queries on the data cube plus the time for creating the data cube is faster than the execution of the example queries over the base data. As we are going to show with different workloads of example queries in our evaluation, we are able to achieve an average speedup of 63. We also compare our approach to standard query optimization with index structures on the base data and show their limited benefit for this use case.

Contributions in Detail and Outline. Our *aggregate-based training phase* consists of two phases: (i) creation of a set of meaningful *pre-aggregated data sets* using data cubes and (ii) rewrite and execute the example queries on the corresponding data cubes or the base data. In detail, the contributions in this paper are:

1. We start with a general overview of ML processes in DBMS in Section 2. In particular, we detail cardinality estimation as a case study for ML in DBMS. We introduce *global* and *local models* as two representatives for ML-based cardinality estimation approaches. Primarily, we show their properties in terms of example workload complexity and conclude the need for optimization of such workloads.
2. Based on this discussion, we introduce our general solution approach of an *aggregated-based training phase* by *pre-aggregating* the base data using the *data cube* concept and

executing the example queries over this pre-aggregated data. Moreover, we introduce a *benefit criterion* to decide whether the *pre-aggregation* is beneficial or not.

3. In Section 4, we present our *aggregate-based training phase* for ML-based cardinality estimation approaches in detail. Our approach consists of two components: *Analyzer* and *Rewrite*. While the main task of the *Analyzer* component is to find and build all beneficial data cubes, the *Rewrite* component is responsible for rewriting the example queries to the constructed data cubes if possible.
4. Then, we present experimental evaluation results for four different workloads for the training phase of ML-based cardinality estimation in Section 5. The workloads are derived from different ML-based cardinality estimation approaches [Ki19b, Wo19b] on the IMDB data set [IM17]. Moreover, we compare our approach with the optimization using index structures.

Finally, we conclude the paper with related work in Section 6 before concluding in Section 7.

2 Machine Learning Models for DBMS

In this section, we start with a brief description of the general process of *machine learning (ML)* in the context of DBMS. Moreover, we discuss ML-based cardinality estimation for DBMS as an important case study and revisit two ML-based approaches solving this specific challenge. Finally, we analyze the specific query workloads for the training phases and clearly state the need for optimized database support.

2.1 Machine Learning Support for DBMS

Most ML-supported techniques for DBMS are supervised learning problems. In this category, there are amongst others: cardinality estimation [Ki19b, Li15, Wo19b], plan cost modeling [MP19, SL19], and indexing [Kr18]. The proposed ML solutions for those highly relevant DBMS problems have a general process in common as shown in Figure 1. This process is usually split into two parts: *forward pass* and *training phase*.

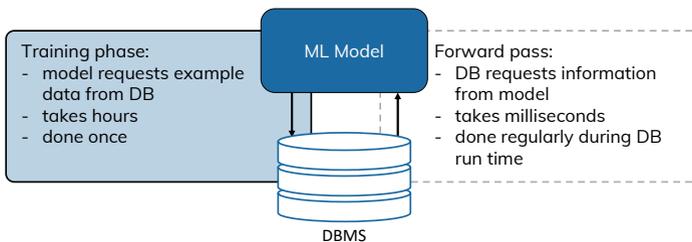


Fig. 1: The general process of supervised ML in DBMS.

Forward pass: This pass consists of the application of the model triggered by a request of the DBMS to the ML model. Each request pulls some specific information from the model such as an estimated cardinality or an index position [Ki19b, Kr18, Li15, Wo19b]. The execution time of each forward pass request is normally in the range of milliseconds. This is advantageous because forward passes occur often and regularly during the run time of the DBMS [Ki19b, Kr18, Li15, Wo19b].

Training phase: To enable the forward pass, a training phase is necessary to construct the required ML model, whereby the challenge for the model lies in the generalization from example data [Ki19b, Kr18, Li15, Wo19b]. Therefore, the model usually requests a lot of diverse labeled example data—pairs of (input, output)—from the DBMS to learn the rules of the underlying problem. Even though the training is performed once, its run time may take hours. This is mainly caused by the generation and execution of a large number of queries against the DBMS to determine the correct cardinalities.

As a consequence, the *training phase* of ML models to support DBMS usually generates a spike high load on the DBMS. Compared to the *forward pass*, the training is significantly more expensive from a database perspective. Therefore, the training phase is a good candidate for optimization to reduce (i) the time for the training phase and (ii) the spike load on the DBMS. Thus, database support or optimization of the training phase is a novel and interesting research field leading to an increased applicability of ML support for DBMS.

2.2 Case Study: Cardinality Estimation

As already mentioned in the introduction, we restrict our focus to the ML-based cardinality estimation use case [Ki19b, Li15, Wo19b]. In this setting, each *forward pass* requests an estimated cardinality for a given query from the ML model. In the *training phase*, the ML cardinality estimator model requires example queries as example data from the DB where the queries are labeled with their true cardinality resulting in pairs of (query, cardinality). These cardinalities are retrieved from the DB by executing the queries enhanced with a count aggregate. Two major approaches for user-workload-independent cardinality estimation with ML models have been proposed in recent years: *global* and *local models*.

Global Model Approach

A *global model* is trained on the *complete database schema* as the underlying model context [Ki19b]. It is effective in covering correlations in attributes for high-quality estimates [Ki19b]. In Figure 2, this is depicted by a single model stated and mapped to the complete schema. Global models have downsides like (i) the complexity of the ML model and (ii) the very expensive training phase. Both disadvantages arise for the following reason: the single ML model handles all attributes and joins in the same way leading to a huge problem space. This huge problem space is directly translated to the model complexity as

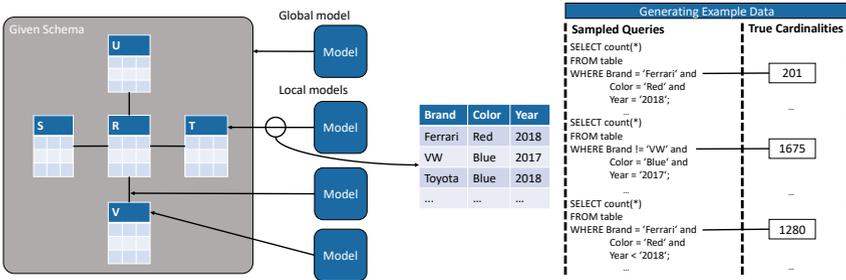


Fig. 2: Overview of ML-based cardinality estimation approaches.

well as to a high number of example queries to cover all predicates and joins over the whole schema as shown [Ki19b].

Local Model Approach

To overcome the shortcomings of the global model approach, the concept of *local models* has been introduced [Wo19b]. Local models are ML models which only cover a certain sub-part of the complete database schema as their model context. This can be a base table or any n-way join. Again, Figure 2 details several local models each covering a different part of the schema. As each of them focuses on a part of the schema, there are many advantages compared to global models. Firstly, local models produce estimates of the same quality as global models [Wo19b]. Secondly, their model complexity is much smaller, because they cover a smaller problem space in different combinations of predicates and joins. The lower complexity stems from a more focused or localized problem solving. A local model has to generalize a smaller problem than the global, i.e. the cardinality estimate of a sub-part of a schema and not the whole schema at once. The lower complexity leads to faster example query sampling and training because the easier problem requires fewer queries during training. A major disadvantage of local models is the high number of models needed to cover all objects touched by a query within the forward pass. Therefore, a separate local ML model needs to be constructed for each part of the schema. Additional queries need to be generated because every local model requests the same amount of example queries. However, these queries are less complex because there are fewer combinations of predicates and tables in a local context.

2.3 Training Phase Workload Analysis

Fundamentally, the global as well as the local ML-based cardinality estimation approach use the same method to sample example queries for the *training phase*. This procedure is shown on the right side of Figure 2, where a local model is trained on an example table

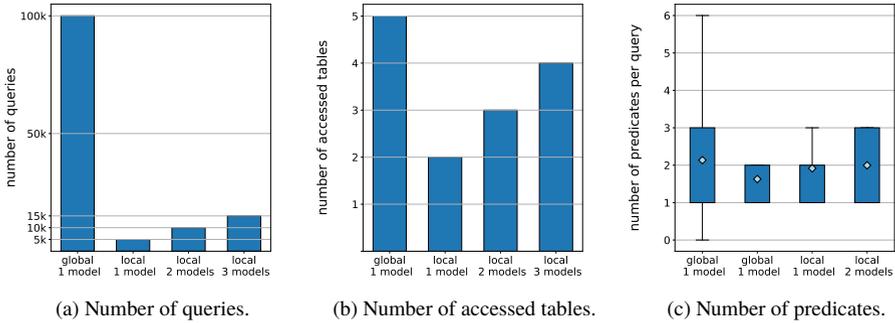


Fig. 3: Analysis of ML-workload complexity for 1 global and 1 to 3 local models for the IMDB.

consisting of three attributes Brand, Color, and Year. To train the local *data-dependent* ML-model, a collection of count queries with different predicate combinations over the table is generated. In our example, the predicates are specified over the three attributes using different operators \neq , $<$, \leq , $=$, $>$, \geq and different predicate values. Thus, all queries have the same structure but differ in their predicates to cover every aspect of the data properties in the underlying table. An example query workload to train a global model would look similar. However, the different contexts for global and local models have an impact on the number and the complexity of the example queries. In general, the query complexity is given by the combinations of joined tables and predicates in a query. The larger the model context, the more complex the example queries. Thus, global models have workloads with a higher number of variations for predicates and joins because they cover the whole schema. Local models are trained with workloads with lesser variation [Wo19b].

To better understand the query workload complexity for the training phase, we analyzed the workloads published by the authors of the global [Ki19a] and local approach [Wo19a]. In both cases, the authors used the IMDB database [IM17] for their evaluation, because this database contains many correlations and is thus very challenging for cardinality estimation. Our analysis results are summarized in Figure 3 and 4.

For a global model and an increasing number of local models, Figure 3a shows the workload complexity in terms of numbers of example queries used per workload. While the global model requires up to 100,000 example queries for the IMDB database [Ki19b], the local model only requires 5,000 example queries per local model [Wo19b] to determine a *stable data-dependent* ML model. In general, the number of example queries is much higher for a global model, but the number of example queries also increases with the number of local models. Thus, we can afford more local models before their collective query count exceeds the number of queries for the global model.

Figure 3b specifies the workload size in terms of data access through the number of joined tables per workload. Similar to the previous figure, the global model has the highest complexity because it requires example queries over more tables to cover the whole schema

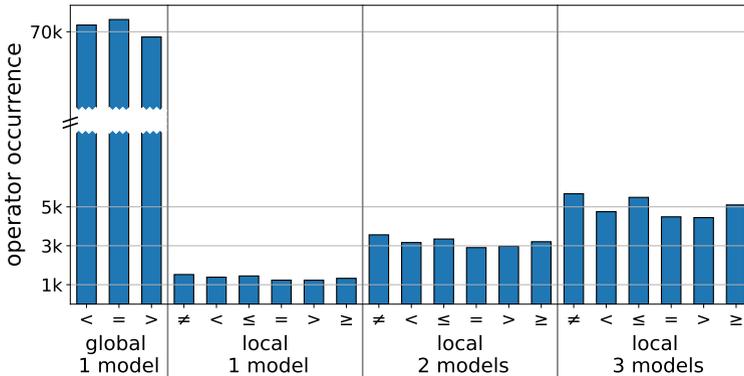


Fig. 4: Predicate operator occurrences for IMDB.

at once. Each local model covers only a limited part of the schema and therefore queries fewer tables per model. Thereby, the complexity of accessed data for local models increases with the number of models.

Another important aspect to describe the workload complexity is the number of predicates per query as shown in Figure 3c. Here, the global model workload has a much larger spread over the number of predicates. The local models detail a more focused distribution with little variation. Again, the local model workload does not require the amount of alternation in predicates of a global model workload because it covers a smaller fragment of the schema. Additionally, Figure 4 gives an overview of the distribution of occurrences of all predicate operators in the workloads as a box plot with mean values. The global model workload only uses the operators $<$, $=$, $>$, whereas the local model workloads use the full set of operators \neq , $<$, \leq , $=$, $>$, \geq . As described by both authors, the predicate operators in each example query are sampled from a uniform distribution [Ki19b, Wo19b]. The slight variation between the operators per workload is due to the fact that both approaches filter 0-tuple queries which do not occur uniformly.

3 Training on Pre-Aggregated Data

As discussed above, the global as well as the local ML-based approach for cardinality estimation generates many example queries with a count aggregate function during the *training phase*. Depending on the model context, there is a small variance in the number of accessed tables, but there is a high variance for predicates in terms of (i) number of predicates, (ii) used predicate operators, and (iii) predicate values in general. So, many queries work on the same data but look at different aspects. Executing such workloads in a naïve way, i.e. executing each example query individually on large base data, is very expensive and generates a high spike load on the database system. The utilization of index

structures for an optimized execution in database systems appears to be an ideal technique at a first glance. However, their benefit is limited as we will show in our evaluation. The same can be said about materialized views. We omitted their evaluation because the experiment did not finish within 30 days.

To tackle this problem more systematically, our core idea is to *pre-aggregate* the base data for different predicate combinations and to reuse this *pre-aggregated* data for several example queries. In general, aggregates compress the data by summarizing information and reducing redundancy. This lessens the amount of data to be scanned by each example query because the aggregates can be smaller than the original data. The aggregate pre-calculates information with the result that the workload queries need to scan less data during execution. Therefore, it is important that the construction of the aggregate does not take longer than the reduction of the workload execution time.

3.1 Grouping Sets as Pre-aggregates

It might sound expensive to aggregate all possible combinations of predicates. However, DBMS already offer substantial supportive data structures for this kind of aggregation. The basic idea of such grouping comes from *Online Analytical Processing (OLAP)* workloads. These aggregate-heavy workloads spawned the idea of pre-aggregating information in *data cubes* [Gr96] helping to reduce the execution time of OLAP queries by collecting and compressing the large amount of data accessed into an aggregate. The concept of data cubes is well-known from data warehouses for storing and computing aggregate information and almost all database systems are offering efficient support for data cube operations [Ag96, Gr96, HRU96, Sh96, ZDN97].

Each attribute of a table or join generates a dimension in the data cube and the distinct attribute values are the dimension values. The cells of a data cube are called facts and contain the aggregate for a particular combination of attribute values. To instantiate the concept of a data cube in a DB, there are different *cube operators*. Usually, these are CUBE, ROLLUP, and GROUPING SET. The CUBE operator instantiates aggregates for the power set of combinations of attribute values. The ROLLUP operator builds the linear hierarchy of

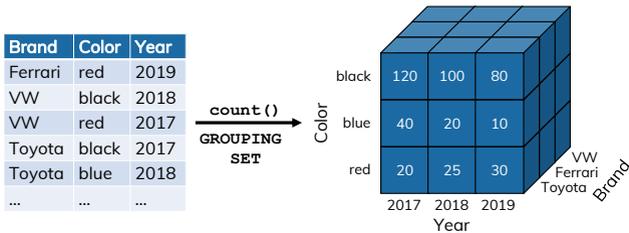


Fig. 5: Aggregating information with grouping sets.

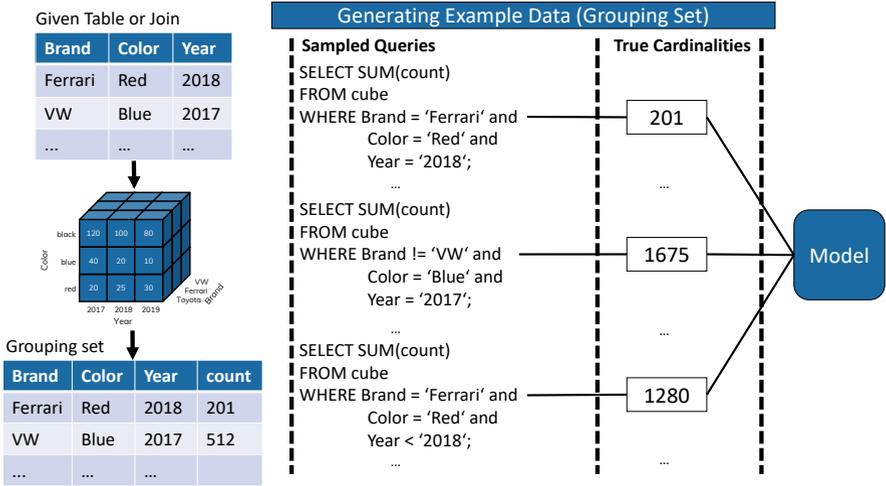


Fig. 6: Illustration of database-supported training phase based on pre-aggregated data/grouping sets.

attribute combinations. The GROUPING SET operator only constructs combinations with all attributes. This characteristic of a grouping set is the major advantage for our use case. With the grouping set aggregation, we compress the original data and avoid the calculation of unnecessary attribute combinations. Figure 5 details an example of a grouping set for a count aggregate over discrete data. The example data has a multidimensional structure after aggregating where each dimension is a property of a car. The cells of the grouping set are filled with the aggregate value, i.e. the count of cars with a particular set of properties.

Given the grouping set data structure, we adapt the generation of example queries from Figure 2. By introducing the data cube, we add an intermediate step before executing the workload. This step constructs a data cube, i.e. grouping set, and rewrites the workload to fit the grouping set. Figure 6 details the construction and the rewrite of queries for the sample data. On the left side, the construction builds a table matching the multidimensional character of the grouping set. Due to this new table layout, the rewrite must include a different aggregate function as shown in Figure 6. For a count aggregate, the corresponding function is a sum over the pre-aggregated count. Last, the rewritten workload is executed and retrieves the output-cardinalities. After the sampling of example queries, the queries and cardinalities are fed to the ML model in the same way as in the original process. This is an advantage of our approach because it does not interfere with other parts of the training process. Therefore, it is independent of the type of ML model and can be applied to a multitude of supervised learning problems.

Even though our approach is independent of the ML model, it is not independent of the data. In general, grouping sets are only beneficial if the aggregate is smaller than the original data. A negative example would be an aggregate over several key columns. Here, the number

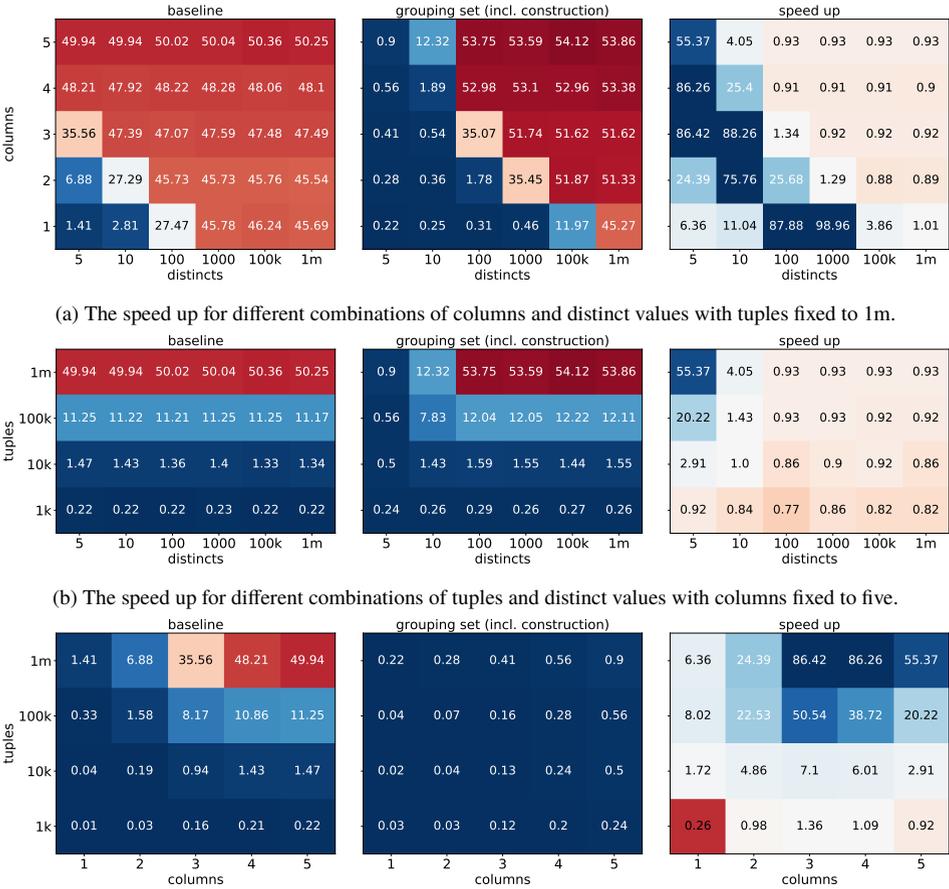


Fig. 7: Execution times and speed ups on synthetic data.

of distinct values per column equals the number of tuples in the table. If such a grouping set is instantiated, each of its dimensions has a length equal to the number of tuples. This grouping set is larger than the original data because it includes the aggregate. With all that in mind, we need to quantify the benefit of a grouping set for our use cases.

3.2 Benefit Criterion

To find a useful criterion when to instantiate grouping sets, we evaluate the usefulness of these sets on synthetic data. Again, we need to find a way to express the compression of information in a grouping set. From the synthetic data, we will derive a general rule for

the theoretical improvement of a grouping set on a given table or join. Our experiment comprises six steps per iteration. The first step generates a synthetic table given three properties. These properties are: the number of tuples in a single table or in the largest table of a join N , the number of columns C , and the number of distinct values per column D . We vary the properties in the ranges:

$$N \in \{1\,000, 10\,000, 100\,000, 1\,000\,000\} \quad (1)$$

$$C \in \{1, 2, 3, 4, 5\} \quad (2)$$

$$D \in \{5, 10, 100, 1\,000, 100\,000, 1\,000\,000\} \quad (3)$$

Changing one property per iteration, this leads to $|N| \cdot |C| \cdot |D| = 4 \cdot 5 \cdot 6 = 120$ different tables or iterations. The values in a column are uniformly sampled from the range of distinct values. With an increasing number of distinct values per column, we simulate floating point columns which have a large number of different values. Columns with few distinct values resemble categorical data. In the second step, we sample 1,000 count aggregate queries as an example workload over all possible combinations of columns (predicates), operators, and values in this iteration. In the third step, we execute these queries against the table and measure their execution time. This is equivalent to the standard procedure to sample example cardinality queries for an ML model. The fourth step constructs the grouping sets over the whole synthetic table and measures its construction time. Next, in step five, we rewrite the queries in a way that they can be executed against the grouping set. We measure their execution time on this grouping set. In the final step, we divide the execution time of the workload on the grouping set by the run time of the workload on the table. This *speed up factor* ranges from close to zero for a negative speed up to infinity for a positive speed up. All time measurements are done three times and averaged. We use PostgreSQL 10 for the necessary data management.

Figure 7 shows the results of all 120 iterations. Blue values mean either better execution times or higher speed up, whereas red means longer execution times or lower speed up. White indicates a speed up factor of one. The first column shows the execution time of the workload against the table. The next column shows the execution time of the workload against the grouping set including the construction time of the grouping set. The last column is the quotient of the second and first column. This is the achieved speed up by using a grouping set. In each row, only two properties are changed while the third property is kept fixed. The first row keeps the number of tuples, the second row the number of columns, and the last row the number of distinct values fixed. From this figure, we can derive three conclusions.

First, we notice that few distinct values in a few columns are beneficial for the aggregation. Next, the more tuples N are in a table, the more distinct values per column can be there for the speed up to be sustained. As the last conclusion, this also applies to the number of columns. All in all, the larger the original table the more distinct values and columns still lead to a speed up. Our experiments show that a grouping set is only beneficial if its size is smaller than the original table. Only then the aggregate compresses information and causes

```

Input  : workload  $wl$ , database
Output : grouping sets

1 grouping sets: tables  $\rightarrow$  attributes
2 for query  $q$  in  $wl$  do                                     // Analyze
3   | tables, attributes of  $q$                                  // step 1
4   | grouping sets[tables] =
5   |   grouping sets[tables]  $\cup$  attributes
6 end

7 for grouping set  $gs$  in grouping sets do
8   |  $N = \max$  (|tuples| of all tables of  $gs$ )                // step 2
9   | for attribute  $p$  in attributes of  $gs$  do
10  |   |  $dv = dv \cup$  |distinct values of  $p$ |
11  |   end
12  |   scaling factor =  $\frac{\prod dv}{N}$                                // step 3
13  |   if scaling factor  $\geq 1$  then
14  |     | grouping sets[tables] = attributes where  $\frac{\prod dv}{N} < 1$ 
15  |     | construct grouping set
16 end

```

Algorithm 1: Analyze component.

less data to be scanned by the queries. Given our evaluation, this happens if the product of the distinct values of all columns is smaller than the table size. We can model this as an equation to be used as a criterion for instantiating beneficial grouping sets.

$$\text{scaling factor} = \frac{1}{N} \prod_{c=1}^C |\text{distinct_values}(\text{column}_c)| \quad (4)$$

If this *scaling factor* is smaller than one, we call a grouping set beneficial. The scaling factor is also a measure of data compression. Therefore, it shows how much faster the scan over the aggregated data can be.

4 Implementation

In this section, we describe the implementation of our *aggregate-based training phase* for ML-based cardinality estimation in DBMS in detail. In our implementation, we assume that a regular DBMS with an SQL interface provides the base data and the ML models are trained outside the DBMS, e.g., in Python. Based on this setting, we added a new layer implemented in Python² between these systems to realize our *aggregate-based training phase* in a very flexible way. Thus, the input of this layer is an ML workload that is necessary for training the ML model. Then, the main tasks of this layer are:

² We plan to make the code base publicly available in case of paper acceptance.

Input : workload wl , grouping sets

Output : rewritten workload wl'

```

1 for query q in wl do                                     // Rewrite
2   tables, attributes of q
3   if attributes = grouping sets[tables] then
4     | rewrite q to match grouping set
5     | add q to wl'
6 end

```

Algorithm 2: Rewrite component.

1. discover as well as create as many beneficial grouping sets in the DBMS as possible for the given ML workload
2. rewrite as well as execute the workload queries according to the grouping sets and base data.

The output of this layer is an annotated ML workload with the retrieved cardinalities on which the ML model is trained afterward. To achieve that, our layer consists of two components. The first component is the *Analyzer* which is responsible for the construction of beneficial grouping sets. The second component is the *Rewrite* rewriting and executing the queries of the ML workload against the constructed grouping sets. In the following, we introduce both components in more detail.

4.1 Analyzer Component

Algorithm 1 gives a more detailed overview over the *Analyzer* Component. Given an ML workload and a database instance, our *Analyzer* consists of three steps to find and build all beneficial grouping sets. In **step one**, the *Analyzer* scans all queries in the ML workload and collects all joins or tables and their respective predicates in use. This generates all possible grouping sets as a mapping from tables building the grouping set to the predicates on those tables. In Algorithm 1, this is covered in lines 1 to 6. The **second step** then collects the number of distinct values per predicate attribute (regardless of their type) and the maximum number of tuples of all tables in the grouping set from the metadata (statistics) of the database. This is shown in lines 8 to 11. In the **third and final step**, our defined benefit criterion (Equation (4)) is used to calculate the scaling factor and therefore the benefit of each grouping set. If the scaling factor is smaller than one, the *Analyzer* constructs the grouping set with all collected predicates. If the scaling factor is larger than or equal to one, the grouping set is constructed with the maximum number of predicates where the scaling factor still is smaller than one. This may disregard certain queries that are subsequently not executed against the grouping set if they have more predicates than the grouping set. On the other hand, queries on the table or join with the predicates in the grouping set can still benefit from it. Moreover, all queries to be executed against a grouping set are marked for rewriting. This final step is detailed in lines 12 to 15.

4.2 Rewrite Component

With all beneficial grouping sets instantiated by the Analyzer Component, it is necessary to modify the ML workload queries to be able to use the pre-aggregates. For this, all queries which can be run against any grouping set will be rewritten in the Rewrite component. The Rewrite component receives information about each query from the Analyzer and rewrites queries in a way that they can be executed against the grouping sets. All queries where the Analyzer does not recognize a grouping set are kept as they are and will be executed over the base data. The Rewrite component is described in Algorithm 2.

When all queries have been processed, the optimized workload is executed as a whole on the database. If a query has been rewritten, it will be executed against the grouping set, otherwise, it will be executed against the original data. Finally, the retrieved results (i.e. cardinalities) are forwarded to the ML system to train the ML model.

5 Evaluation

To show the benefit of our novel *aggregate-based training phase*, we conducted an exhaustive experimental study with both presented types of ML models for cardinality estimation (cf. Section 2). Thus, we start this section by explaining the experimental settings followed by a description of selective results for the local as well as global ML model approaches. Afterward, we summarize the main experimental findings.

5.1 Experimental Setting

For our experiments, we used the original workloads for the local and global ML model approaches [Ki19a, Wo19a] on the IMDB data set [IM17]. The IMDB contains a snowflake database schema with several millions of tuples in both the fact and the 20 dimension tables. As already presented in Section 2.3, the global model workload contains 100,000 queries. For the local models, we used three workloads where each workload has 5,000 queries more

model	Base Data w/o Index	Base Data w/ Index	Construction GS	Execution GS	Total GS	Coverage GS
local 1	2h 30m	1h 44m	6.17s	191.34s	197.51s	100%
local 2	7h 53m	5h 10m	23.70s	205.91s	229.61s	100%
local 3	10h 12m	6h 56m	29.10s	430.36s	459.46s	100%
global full	–	4d 14h	2h 22m	20d 20h	20d 22h	100%
global opt	–	4d 14h	34m 29s	2d 11h	2d 12h	55%

Tab. 1: Execution times ML workloads (GS: grouping sets).

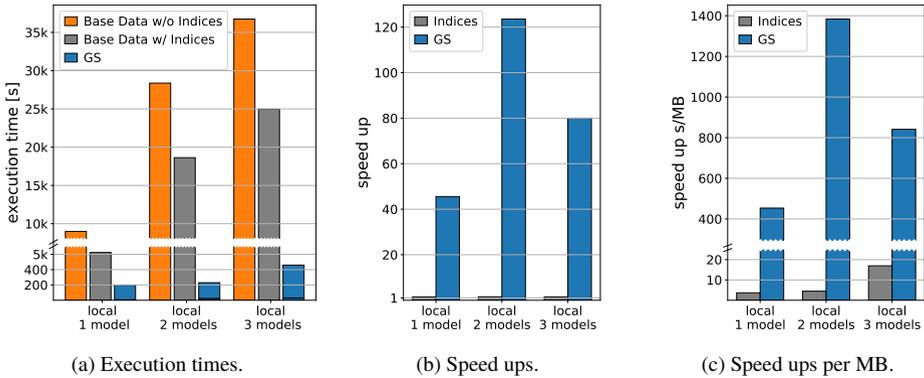


Fig. 8: Local model evaluation based on our *aggregate-based training phase* (GS: Grouping Sets).

than the previous one. These workloads correspond to one, two, and three trained local models. Overall, we have four workloads for our experiments: one for a global model and three workloads for an increasing number of local models. Moreover, all experiments are conducted on an AMD A10-7870K system with 32GB main-memory with PostgreSQL 10 as the underlying database system.

In our experiments, we measured the workload execution times, whereby we distinguish three different execution modes:

Base Data w/o Indexes: ML workload is executed on the IMDB base data without any indexes on the base data.

Base Data w/ Indexes: ML workload is executed on the IMDB base data with indexes on all (join) predicates in use.

Grouping Set (GS): ML workload is executed on *pre-aggregated data* as determined by our approach.

The first two execution modes represent our baselines because both are currently used in the presented ML model approaches for cardinality estimation [Ki19b, Wo19b, Li15].

5.2 Experimental Results: Local Model Workloads

Figure 8 shows the results for the three local model workloads. The first workload contains the necessary queries to build one local ML model to estimate the cardinalities for the join `title>movie_keyword`. The second workload adds 5,000 queries to the first workload to construct a second ML model for the join `title>movie_info`. The third workload adds another ML model for an additional join `title>movie_companies`. Therefore, we increment the number of local ML models.

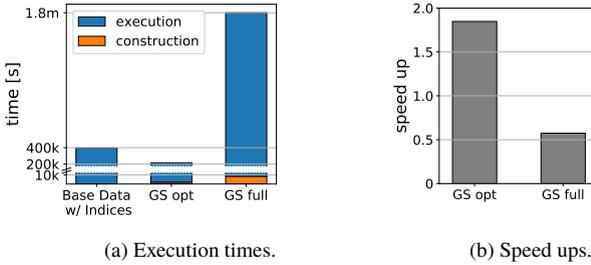


Fig. 9: Global model evaluation results.

Figure 8a details the execution times for all three local model workloads for all investigated execution modes. In each of the three groups, the left bar shows the complete workload execution time on the IMDB without indexes, the middle bar on the IMDB with indexes, and the right bar the execution time with our grouping set approach. As we can see, indexes on the base data are already helping to reduce the workload execution times compared to execution on the base data w/o indexes, but the speedup is very marginal as shown in Figure 8b because the DBMS might decide to abstain from using the indexes. In contrast to that, our grouping set approach has the lowest execution times in all cases and the achieved speed ups compared to execution on the base data w/o indexes are in the range between 45 and 125 as depicted in Figure 8b. Thus, we can conclude that our *aggregation-based training phase* is much more efficient than state-of-the-art approaches.

For each considered join, our *aggregation-based* approach creates a specific grouping set containing all columns from the corresponding workload queries. According to equation (4), the scaling factors are: 0.02, 0.003, and 0.05 for the three joins. Thus, the instantiation of grouping sets is beneficial. So, all grouping sets achieve a very good compression rate and the rewritten workload queries on the grouping sets have to read much less data compared to the execution on base data. Moreover, all workload queries can be rewritten, so that the coverage is 100% and every query benefits from this optimization. Nevertheless, the three scaling factors differ explaining the different speed ups.

The construction of the grouping sets can be considered a drawback. However, as illustrated in Table 1, the construction times for the grouping sets are negligible because the reduction in execution time is significantly higher. From a storage perspective, index structures and grouping sets need some extra storage space, where the storage overhead for grouping sets is larger than for indices. But, as illustrated in Figure 8c, the speed up per additional MB for grouping sets is much larger than for indices. All in all, we gain a much larger speed up making grouping sets the more efficient approach.

5.3 Experimental Results: Global Model Workload

Figure 9 shows the evaluation results in terms of execution and construction times for the global model workload. As shown in the previous experiment, the utilization of indices is always beneficial. Thus, we only compare the execution on base data with indices and the execution on the aggregated data in this evaluation.

In general, there are 21 grouping sets possible for the global workload. However, some of these grouping sets have a scaling factor larger than one. Therefore, our Analyzer component disregards the attributes of some grouping sets until the scaling factor is smaller than one (cf. Section 4). As a consequence, only 55% of the global workload queries can be rewritten to this optimal set of grouping sets. This strategy called *GS opt* in Figure 9a reduces the workload execution time of the global workload. The speed up compared to the execution on the base data with indexes is almost 2 (Figure 9b).

To show the benefit of our grouping set selection strategy, we also constructed all grouping sets with all attributes (*GS full*). There, we are able to rewrite all global workload queries to be executed on these aggregated data. As shown in Figure 9a, the overall workload execution time is longer than the execution on base data with indices. Therefore, grouping sets have to be selected carefully. Moreover, this experiment shows that our definition of a beneficial grouping set is applicable because (i) not all grouping sets are beneficial and (ii) not all queries can or need to be optimized with a grouping set. The benefit criterion considers both aspects to reduce workload execution times.

5.4 Main Findings

For both types of ML models for cardinality estimation, our *aggregation-based training phase* offers a database-centric way to reduce execution time. Table 1 summarizes our evaluation results. The overhead introduced by the construction of a grouping set is much smaller than the savings in execution time. So, grouping sets reduce the workload execution times and amortize their own construction time. The simpler structure of the local model workloads is better supported by grouping sets because they contain fewer combinations of columns and fewer distinct values. These are exactly two of the assets for grouping sets identified in Section 3. This leads to a higher performance speed up for local model workloads than for global model workloads with consistent high-quality estimates. Thus, we can afford a larger amount of local models to reach the schema coverage of a global model. Even if these models request more queries than the global model, their benefits from the use of grouping sets outweigh the higher number of queries.

6 Related Work

In this section, we detail the importance of database support for machine learning in other works. We look at the motivation for pre-aggregates from both the database system and the machine learning point of view.

When looking at the synergy of database and machine learning systems, there are three possible interactions: (i) integrate machine learning into database systems, (ii) adapt database techniques for machine learning models, and (iii) combine database and machine learning into one life cycle system [KBY17]. Based on that, we classify our work in category (iii). However, the focus in this area is more on feature and model selection and not on sampling example data. We argue that the direct support of machine learning training phases with databases should be treated with the same attention.

To the best of our knowledge, there is only little research on directly optimizing the sampling of workloads for machine learning problems. The authors of [Ki19b] detail their method of speeding up query sampling in [Ki19c]. They use massive parallelism by distributing the workload over several DB instances. We see this as a promising step because our approach can also profit from parallel execution. Especially the instantiation and the querying of grouping sets can be done in parallel because grouping sets are orthogonal to each other.

Another thing to look at is the availability of supportive data structures in database systems. The cube operators are established in databases and benefit from a wide-ranged support [Ag96, Gr96, HRU96, Sh96, ZDN97]. The ability of a database to deliver necessary meta information is also important. For example, fast querying for the distinct values of each column has a large impact on performance. A simple solution for this is a dictionary encoding of the data in the database. Some database systems already use dictionary coding for all their data [Fä12]. This is beneficial for our approach because from a dictionary encoded column it is easy to yield the number of distinct values with a dictionary scan. Moreover, dictionary encoding directly supports the transformation of the data into the grouping set dimension and the definition of ranges of these dimensions.

Aside from machine learning, database support for data mining has already been an important research topic [AS96, CWC09, HLH03, Ne01, OC00]. For example, [HLH03] identifies that aggregation in sub-spaces formed by combinations of attributes is a common task in many data mining algorithms. Based on that observation, we see a large potential for tighter coupling of databases and mining algorithms.

7 Conclusion

We made the case for cardinality estimation as a candidate for database support of machine learning for DBMS. We detailed an approach for pre-aggregating count information for cardinality estimation workloads. It uses grouping sets, a well-known database operator, to

reduce the data to be scanned by example queries for cardinality estimation with machine learning models. This reduces the execution time of a given workload even though we spend extra time to construct the intermediate data structures.

This case has a strong potential to be applied to the other similar machine learning problems, like plan cost modeling or indexing. We see parallels between the potential for machine learning workloads and any of these machine learning problems where information about the data in the DB is aggregated. These parallels make grouping sets and therefore DB support beneficial for ML for DBMS in general.

Bibliography

- [Ag96] Agarwal, Sameet; Agrawal, Rakesh; Deshpande, Prasad; Gupta, Ashish; Naughton, Jeffrey F.; Ramakrishnan, Raghu; Sarawagi, Sunita: On the Computation of Multidimensional Aggregates. In: VLDB. pp. 506–521, 1996.
- [AS96] Agrawal, Rakesh; Shim, Kyuseok: Developing Tightly-Coupled Data Mining Applications on a Relational Database System. In: KDD. pp. 287–290, 1996.
- [CWC09] Cho, Chung-Wen; Wu, Yi-Hung; Chen, Arbee L. P.: Effective database transformation and efficient support computation for mining sequential patterns. *J. Intell. Inf. Syst.*, 32(1):23–51, 2009.
- [Da19] Damme, Patrick; Ungethüm, Annett; Hildebrandt, Juliana; Habich, Dirk; Lehner, Wolfgang: From a Comprehensive Experimental Survey to a Cost-based Selection Strategy for Lightweight Integer Compression Algorithms. *ACM TODS*, 44(3):9:1–9:46, 2019.
- [Fä12] Färber, Franz; May, Norman; Lehner, Wolfgang; Große, Philipp; Müller, Ingo; Rauhe, Hannes; Dees, Jonathan: The SAP HANA database - An architecture overview. *IEEE Data Eng. Bull.*, 35:28–33, 03 2012.
- [FM11] Fender, Pit; Moerkotte, Guido: A new, highly efficient, and easy to implement top-down join enumeration algorithm. In: ICDE. pp. 864–875, 2011.
- [Gr96] Gray, J.; Bosworth, A.; Lyaman, A.; Pirahesh, H.: Data cube: a relational aggregation operator generalizing GROUP-BY, CROSS-TAB, and SUB-TOTALS. In: ICDE. 1996.
- [HLH03] Hinneburg, Alexander; Lehner, Wolfgang; Habich, Dirk: COMBI-Operator: Database Support for Data Mining Applications. In: VLDB. pp. 429–439, 2003.
- [HN17] Harmouch, Hazar; Naumann, Felix: Cardinality Estimation: An Experimental Survey. *PVLDB*, 11(4):499–512, 2017.
- [HRU96] Harinarayan, Venky; Rajaraman, Anand; Ullman, Jeffrey D.: Implementing Data Cubes Efficiently. In: SIGMOD. pp. 205–216, 1996.
- [IM17] IMDB: Internet Movie Database. <ftp://ftp.fu-berlin.de/pub/misc/movies/database/frozendata/>, Accessed on 2020-06-01.
- [KBY17] Kumar, Arun; Boehm, Matthias; Yang, Jun: Data management in machine learning: Challenges, techniques, and systems. In: SIGMOD. pp. 1717–1722, 2017.

- [KHL17] Karnagel, Tomas; Habich, Dirk; Lehner, Wolfgang: Adaptive Work Placement for Query Processing on Heterogeneous Computing Resources. *PVLDB*, 10(7):733–744, 2017.
- [Ki19a] Kipf, Andreas: Learned Cardinalities in PyTorch. <https://github.com/andreaskipf/learnedcardinalities/>, Accessed on 2020-08-20.
- [Ki19b] Kipf, Andreas; Kipf, Thomas; Radke, Bernhard; Leis, Viktor; Boncz, Peter A.; Kemper, Alfons: Learned Cardinalities: Estimating Correlated Joins with Deep Learning. In: *CIDR*. 2019.
- [Ki19c] Kipf, Andreas; Vorona, Dimitri; Müller, Jonas; Kipf, Thomas; Radke, Bernhard; Leis, Viktor; Boncz, Peter; Neumann, Thomas; Kemper, Alfons: Estimating Cardinalities with Deep Sketches. In: *SIGMOD*. p. 1937–1940, 2019.
- [Kr18] Kraska, Tim; Beutel, Alex; Chi, Ed H; Dean, Jeffrey; Polyzotis, Neoklis: The case for learned index structures. In: *SIGMOD*. pp. 489–504, 2018.
- [Le15] Leis, Viktor; Gubichev, Andrey; Mirchev, Atanas; Boncz, Peter; Kemper, Alfons; Neumann, Thomas: How good are query optimizers, really? *PVLDB*, 9(3):204–215, 2015.
- [Li15] Liu, Henry; Xu, Mingbin; Yu, Ziting; Corvinelli, Vincent; Zuzarte, Calisto: Cardinality Estimation Using Neural Networks. In: *CASCON*. pp. 53–59, 2015.
- [MNS09] Moerkotte, Guido; Neumann, Thomas; Steidl, Gabriele: Preventing Bad Plans by Bounding the Impact of Cardinality Estimation Errors. *PVLDB*, 2(1):982–993, 2009.
- [MP19] Marcus, Ryan; Papaemmanouil, Olga: Plan-Structured Deep Neural Network Models for Query Performance Prediction. *PVLDB*, 12(11):1733–1746, 2019.
- [Ne01] Netz, Amir; Chaudhuri, Surajit; Fayyad, Usama M.; Bernhardt, Jeff: Integrating Data Mining with SQL Databases: OLE DB for Data Mining. In: *ICDE*. pp. 379–387, 2001.
- [OC00] Ordonez, Carlos; Cereghini, Paul: SQLEM: Fast Clustering in SQL using the EM Algorithm. In: *SIGMOD*. pp. 559–570, 2000.
- [Ro15] Rosenfeld, Viktor; Heimerl, Max; Viebig, Christoph; Markl, Volker: The Operator Variant Selection Problem on Heterogeneous Hardware. In: *ADMS@VLDB*. pp. 1–12, 2015.
- [Sh96] Shukla, Amit; Deshpande, Prasad; Naughton, Jeffrey F; Ramasamy, Karthikeyan: Storage estimation for multidimensional aggregates in the presence of hierarchies. In: *VLDB*. pp. 522–531, 1996.
- [SL19] Sun, Ji; Li, Guoliang: An End-to-End Learning-Based Cost Estimator. *PVLDB*, 13(3):307–319, 2019.
- [Wo19a] Woltmann, Lucas: Cardinality Estimation with Local Deep Learning Models. <https://github.com/lucaswo/cardest/>, Accessed on 2020-08-20.
- [Wo19b] Woltmann, Lucas; Hartmann, Claudio; Thiele, Maik; Habich, Dirk; Lehner, Wolfgang: Cardinality Estimation with Local Deep Learning Models. In: *aiDM*. ACM, 2019.
- [YW79] Youssefi, Karel; Wong, Eugene: Query Processing in a Relational Database Management System. In: *VLDB*. pp. 409–417, 1979.
- [ZDN97] Zhao, Yihong; Deshpande, Prasad; Naughton, Jeffrey F.: An Array-Based Algorithm for Simultaneous Multidimensional Aggregates. In: *SIGMOD*. pp. 159–170, 1997.