Aggregate-based Training Phase for ML-based Cardinality Estimation

Lucas Woltmann¹ · Claudio Hartmann¹ · Dirk Habich¹ · Wolfgang Lehner¹

Received: 27 October 2021 / Accepted: 13 December 2021 / Published online: 10 January 2022
© The Author(s) 2022

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 90 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 [1]. 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 [1–4]. Based on these estimations, various decisions are made by different optimization techniques such as choosing (i) the right join order [5], (ii) the right physical operator variant [6], (iii) the best-fitting compression scheme [7], or (iv) the optimal operator

placement within heterogeneous hardware [8]. However, to make good decisions in all cases, it is important to have cardinality estimations with high accuracy.

As shown in recent papers [9, 10], including our own work [11], 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 [1], ML-based approaches assume that a sufficiently deep neural network can model the very complex data dependencies and correlations [9]. For this reason, ML-based cardinality estimation approaches may thus give much more accurate estimations as clearly demonstrated in [9–11]. 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 [9, 11], whereby the execution of those example queries can be very time consuming, especially for

✉ Lucas Woltmann
lucas.woltmann@tu-dresden.de

Claudio Hartmann
claudio.hartmann@tu-dresden.de

Dirk Habich
dirk.habich@tu-dresden.de

Wolfgang Lehner
wolfgang.lehner@tu-dresden.de

¹ Dresden Database Research Group, Technische Universität Dresden, Dresden, Germany

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. This is an extended version of previous work [12]. Here, we expand the evaluation and introduce further improvement of the *aggregate-based training phase*. Usually, as described in [9, 11], 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 [13]. 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 speed up of 90. 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 Sect. 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 Sect. 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 Sect. 5. The workloads are derived from different ML-based cardinality estimation approaches [9, 11] on the IMDb data set [14]. Moreover, we improve our approach by using additional indexes on the pre-aggregates. We compare our results with the optimization using just index structures on the base data.

Finally, we conclude the paper with related work in Sect. 6 before concluding in Sect. 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 [9–11], plan cost modeling [15, 16], and indexing [17]. The proposed ML solutions for those highly relevant DBMS problems have a general process in common as shown in Fig. 1. This process is usually split into two parts: *forward pass* and *training phase*.

Forward pass: This pass consists of the application of the model triggered by a request of the DBMS to the ML

Fig. 1 The general process of supervised ML in DBMS

model. Each request pulls some specific information from the model such as an estimated cardinality or an index position [9–11, 17]. 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 [9–11, 17].

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 [9–11, 17]. 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 ground truth.

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 [9–11]. 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*.

2.2.1 Global Model Approach

A *global model* is trained on the *complete database schema* as the underlying model context [9]. It is effective in covering correlations in attributes for high-quality estimates [9]. In Fig. 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 well as to a high number of example queries to cover all predicates and joins over the whole schema as shown [9].

2.2.2 Local Model Approach

To overcome the shortcomings of the global model approach, the concept of *local models* has been introduced [11]. 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, Fig. 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 [11]. Sec-

Fig. 2 Overview of ML-based cardinality estimation approaches

only, 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 Fig. 2, where a local model is trained on an example table 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 [11].

To better understand the query workload complexity for the training phase, we analyzed the workloads published by the authors of the global [18] and local approach [19]. In both cases, the authors used the IMDb database [14] for their evaluation, because this database contains many correlations and is thus very challenging for cardinality estimation. Our analysis results are summarized in Fig. 3 and 4.

For a global model and an increasing number of local models, Fig. 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 [9], the local model only requires 5000 example queries per local model [11] 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.

Fig. 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 at once. Each local model covers only a limited part of the schema and therefore queries

Fig. 3 Analysis of ML-workload complexity for 1 global and 1 to 3 local models for the IMDb. **a** Number of queries, **b** number of accessed tables, **c** number of predicates

Fig. 4 Predicate operator occurrences for IMDb

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 Fig. 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, Fig. 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 [9, 11]. 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 naive 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* [13] 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 [13, 20–23].

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 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. Fig. 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 Fig. 2. By introducing the data cube, we add an intermediate step before executing the

Fig. 5 Aggregating information with grouping sets

workload. This step constructs a data cube, i.e. grouping set, and rewrites the workload to fit the grouping set. Fig. 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 Fig. 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 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:

$$N \in \{1000; 10,000; 100,000; 1,000,000\} \quad (1)$$

$$V_{\alpha'} = 1 - e^{-k \cdot (M_S - T_Q)} \quad (2)$$

$$D \in \{5; 10; 100; 1000; 100,000; 1,000,000\} \quad (3)$$

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 1000 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

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

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.

Fig. 7 shows the results of all 120 iterations. A darker cell hue means either better execution times or higher speed up, whereas lighter areas mean longer execution times or lower speed up. 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

Fig. 7 Execution times and speed ups on synthetic data. **a** The speed up for different combinations of columns and distinct values with tuples fixed to 1 m, **b** the speed up for different combinations of tuples and distinct values with columns fixed to five, **c** the speed up for different combinations of tuples and columns with distinct values fixed to five

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 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.

Algorithm 1 Analyze component.

Require: workload wl , database

```

1: for query  $q \in wl$  do ▷ step 1
2:   tables, attributes of  $q$ 
3:   grouping sets[tables] =
4:     grouping sets[tables]  $\cup$  attributes
5: end for
6:
7: for  $gs \in$  grouping sets do ▷ step 2
8:    $N = \max(|tuples| \text{ in tables of } gs)$ 
9:   for attribute  $p \in$  attributes of  $gs$  do
10:     $dv = dv \cup |\text{distinct values of } p|$ 
11:   end for
12:   scaling factor =  $\frac{\prod dv}{N}$  ▷ step 3
13:   if scaling factor  $\geq 1$  then
14:     grouping sets[tables] =
15:       attributes where  $\frac{\prod dv}{N} < 1$ 
16:   end if
17:   Construct grouping set
18: end for
19: return grouping sets: tables  $\rightarrow$  attributes

```

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 Python 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:

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 5. 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 7 to 11. In the **third and final step**, our defined benefit criterion (Eq. (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 18.

Algorithm 2 Rewrite component.

Require: workload wl , grouping sets

```

1: for query  $q \in wl$  do
2:   tables, attributes  $\leftarrow q$ 
3:   if attributes = grouping sets[tables] then
4:     Rewrite  $q$  to match grouping set
5:   end if
6:   Add  $q$  to  $wl'$ 
7: end for
8: return rewritten workload  $wl'$ 

```

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. Sect. 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 [18, 19] on the IMDb data set [14]. 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 Sect. 2.3, the global model workload contains 100,000 queries. For the local models, we used three workloads where each workload has 5000 queries more 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 four 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, totaling 15 indexes.

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

Grouping Set (GS) w/ Indexes: ML workload is executed on *pre-aggregated data* where indexes were built on all predicates in the GS, totaling 39 indexes.

The first two execution modes represent our baselines because both are currently used in the presented ML model approaches for cardinality estimation [9–11].

5.2 Experimental Results: Local Model Workloads

Fig. 8 shows the results for the three local model workloads. The first workload *local 1* contains the necessary queries to build one local ML model to estimate the cardinalities for the join `title` $\bowtie$ `movie_keyword`. The second workload *local 2* adds 5000 queries to the first workload to construct a second ML model for the join `title` $\bowtie$ `movie_info`. The third workload *local 3* adds another ML model for an additional join `title` $\bowtie$ `movie_companies`. Therefore, we increment the number of local ML models.

Fig. 8a details the execution times for all three local model workloads *local 1*, *local 2*, and *local 3* for all investigated execution modes. In each of the three groups, the

Table 1 Execution times ML workloads (GS: grouping sets)

Model	Base Data w/o Index	Base Data w/ Index	Construction GS	Execution GS	Total GS	Total GS w/ Index	Coverage GS
Local 1	2h 30m	1h 44m	6.17s	191.34s	197.51s	139.47s	100%
Local 2	7h 53m	5h 10m	23.70s	205.91s	229.61s	149.99s	100%
Local 3	10h 12m	6h 56m	29.10s	430.36s	459.46s	295.25s	100%
Global full	–	4d 14h	2h 22m	20d 20h	20d 22h	–	100%
Global opt	–	4d 14h	34m 29s	2d 11h	2d 12h	1d 21h	55%

left bar shows the complete workload execution time on the IMDb without indexes, the second bar on the IMDb with indexes, the third bar the execution time with our grouping set approach, and the last bar the execution time with our grouping sets including indexes. As we can see, indexes on the base data are already reducing the workload execution times compared to execution on the base data w/o indexes, but the speedup is very marginal as shown in Fig. 8b because the DBMS might decide to abstain from using the indexes. In contrast to that, our grouping set approach has lower execution times in all cases and the achieved speedups compared to execution on the base data w/o indexes are in the range between 45 and 125 as depicted in Fig. 8b. Additionally, building indexes on the grouping sets improves the speed up to a range of 65 to 180. 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 Eq. (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 and indexes 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 indexes. But, as illustrated in Fig. 8c, the speed up per additional MB for grouping sets is much larger than for indexes. We also show that the use of indexes on grouping sets is beneficial but grouping sets are more efficient regarding memory consumption. So, the additional storage requirements for grouping sets and indexes are justifiable. All in all, we gain a much larger speed up making grouping sets or the combination of grouping sets and indexes the more efficient approach.

Fig. 8 Local model evaluation based on our *aggregate-based training phase* (GS: Grouping Sets). **a** Execution times, **b** speed ups, **c** speed ups per MB

5.3 Experimental Results: Global Model Workload

Fig. 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 indexes is always beneficial. Thus, we only compare the execution on base data with indexes, the execution on the aggregated data, and the execution on the combined aggregated and indexed 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. Sect. 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* in Fig. 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 (Fig. 9b). Using indexes on the pre-aggregated data brings the speed up to 2.5.

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 Fig. 9a, the overall workload execution time is longer than the execution on base data with indexes. 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. Additionally, the usage of indexes is always beneficial including their construction on the grouping sets. 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 Sect. 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 [24]. 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 [9] detail their method of speeding up query sampling in [25]. 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 [13, 20–23]. 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 [26]. This is beneficial for our approach because from

Fig. 9 Global model evaluation results. **a** Execution times, **b** speed ups

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 [27–31]. For example, [29] 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. These practical benefits are also shown in *PostCENN*, an extended version of PostgreSQL, where we evaluated our approach by integrating local models and the *Aggregate-based Training Phase* [32]. This extended PostgreSQL system gives the users the possibility to use and to compare histograms and NNs as cardinality estimators both in quality and performance. Nevertheless, for the broad applicability of NNs for cardinality estimation in database systems, further research is necessary in the following areas: (i) how to deal with data inserts and updates or (ii) how to reduce the training times even more.

Moreover, our presented approach in this paper has a strong potential to be applied to the other similar machine learning problems as well, like plan cost modeling or indexing. In particular, 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.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are

included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Leis V, Gubichev A, Mirchev A et al (2015) How good are query optimizers, really? Proc VLDB Endow 9(3):204–215
2. Harmouch H, Naumann F (2017) Cardinality estimation: an experimental survey. Proc VLDB Endow 11(4):499–512
3. Moerkotte G, Neumann T, Steidl G (2009) Preventing bad plans by bounding the impact of cardinality estimation errors. Proc VLDB Endow 2(1):982–993
4. Youssefi K, Wong E (1979) Query processing in a relational database management system. In: VLDB, pp 409–417
5. Fender P, Moerkotte G (2011) A new, highly efficient, and easy to implement top-down join enumeration algorithm. In: ICDE, pp 864–875
6. Rosenfeld V, Heimerl M, Viebig C et al (2015) The operator variant selection problem on heterogeneous hardware. In: ADMS@VLDB, pp 1–12
7. Damme P, Ungethüm A, Hildebrandt J et al (2019) From a comprehensive experimental survey to a cost-based selection strategy for lightweight integer compression algorithms. ACM Trans Database Syst 44(3):9:1–9:46. <https://doi.org/10.1145/3323991>
8. Karnagel T, Habich D, Lehner W (2017) Adaptive work placement for query processing on heterogeneous computing resources. Proc VLDB Endow 10(7):733–744
9. Kipf A, Kipf T, Radke B et al (2019) Learned cardinalities: Estimating correlated joins with deep learning. In: CIDR
10. Liu H, Xu M, Yu Z et al (2015) Cardinality estimation using neural networks. In: CASCON, pp 53–59
11. Woltmann L, Hartmann C, Thiele M et al (2019) Cardinality estimation with local deep learning models. In: Bordawekar R, Shmueli O (eds) Proceedings of the Second International Workshop on Exploiting Artificial Intelligence Techniques for Data Management, aiDM at SIGMOD 2019 Amsterdam, The Netherlands, 5 July 2019. vol 10. ACM, pp 5:1–5:8 <https://doi.org/10.1145/3329859.3329875>
12. Woltmann L, Hartmann C, Habich D et al (2021) Aggregate-based training phase for ML-based cardinality estimation. In: Sattler K, Herschel M, Lehner W (eds) Datenbanksysteme für Business, Technologie und Web (BTW 2021), 19. Fachtagung des GI-Fachbereichs “Datenbanken und Informationssysteme” (DBIS) Dresden, Germany, 13.–17. September 2021. Proceedings, LNI, vol P-311. Gesellschaft für Informatik, Bonn, pp 135–154 <https://doi.org/10.18420/btw2021-07>
13. Gray J, Bosworth A, Lyaman A et al (1996) Data cube: a relational aggregation operator generalizing group-by, cross-tab, and sub-totals. In: ICDE
14. IMDb (2017) Imdb: Internet movie database. <ftp://ftp.fu-berlin.de/pub/misc/movies/database/frozendata/>. Accessed 1 Sept 2021
15. Marcus R, Papaemmanouil O (2019) Plan-structured deep neural network models for query performance prediction. Proc VLDB Endow 12(11):1733–1746. <https://doi.org/10.14778/3342263.3342646>
16. Sun J, Li G (2019) An end-to-end learning-based cost estimator. Proc VLDB Endow 13(3):307–319. <https://doi.org/10.14778/3368289.3368296>
17. Kraska T, Beutel A, Chi EH et al (2018) The case for learned index structures. In: SIGMOD, pp 489–504
18. Kipf A (2019) Learned cardinalities in pytorch. <https://github.com/andreaskipf/learnedcardinalities/>. Accessed 20 Oct 2021

19. Woltmann L (2019) Cardinality estimation with local deep learning models. <https://github.com/lucaswo/cardest/>. Accessed 20 Oct 2021
20. Agarwal S, Agrawal R, Deshpande P et al (1996) On the computation of multidimensional aggregates. In: VLDB, pp 506–521
21. Harinarayan V, Rajaraman A, Ullman JD (1996) Implementing data cubes efficiently. In: SIGMOD, pp 205–216
22. Shukla A, Deshpande P, Naughton JF et al (1996) Storage estimation for multidimensional aggregates in the presence of hierarchies. In: VLDB, pp 522–531
23. Zhao Y, Deshpande P, Naughton JF (1997) An array-based algorithm for simultaneous multidimensional aggregates. In: SIGMOD, pp 159–170
24. Kumar A, Boehm M, Yang J (2017) Data management in machine learning: challenges, techniques, and systems. In: SIGMOD, pp 1717–1722
25. Kipf A, Vorona D, Müller J et al (2019) Estimating cardinalities with deep sketches. In: SIGMOD, pp 1937–1940 <https://doi.org/10.1145/3299869.3320218>
26. Färber F, May N, Lehner W et al (2012) The SAP HANA database – an architecture overview. IEEE Data Eng Bull 35:28–33
27. Agrawal R, Shim K (1996) Developing tightly-coupled data mining applications on a relational database system. In: KDD, pp 287–290
28. Cho C, Wu Y, Chen ALP (2009) Effective database transformation and efficient support computation for mining sequential patterns. J Intell Inf Syst 32(1):23–51
29. Hinneburg A, Lehner W, Habich D (2003) Combi-operator: database support for data mining applications. In: VLDB, pp 429–439
30. Netz A, Chaudhuri S, Fayyad UM et al (2001) Integrating data mining with SQL databases: OLE DB for data mining. In: ICDE, pp 379–387
31. Ordonez C, Cereghini P (2000) SQLEM: fast clustering in SQL using the EM algorithm. In: SIGMOD, pp 559–570
32. Woltmann L, Olwig D, Hartmann C et al (2021) PostCENN: postgresql with machine learning models for cardinality estimation. Proc VLDB Endow 14(12):2715–2718