



ALECE: An Attention-based Learned Cardinality Estimator for SPJ Queries on Dynamic Workloads

Pengfei Li
Alibaba Group, China
lpf367135@alibaba-inc.com

Wenqing Wei
Alibaba Group, China
weiwenqing.wwq@alibaba-inc.com

Rong Zhu †
Alibaba Group, China
red.zr@alibaba-inc.com

Bolin Ding †
Alibaba Group, China
bolin.ding@alibaba-inc.com

Jingren Zhou †
Alibaba Group, China
jingren.zhou@alibaba-inc.com

Hua Lu †
Roskilde University, Denmark
luhua@ruc.dk

ABSTRACT

For efficient query processing, DBMS query optimizers have for decades relied on delicate cardinality estimation methods. In this work, we propose an Attention-based LEarned Cardinality Estimator (ALECE for short) for SPJ queries. The core idea is to discover the implicit relationships between queries and underlying dynamic data using attention mechanisms in ALECE’s two modules that are built on top of carefully designed featurizations for data and queries. In particular, from all attributes in the database, the data-encoder module obtains organic and learnable aggregations which implicitly represent correlations among the attributes, whereas the query-analyzer module builds a bridge between the query featurizations and the data aggregations to predict the query’s cardinality. We experimentally evaluate ALECE on multiple dynamic workloads. The results show that ALECE enables PostgreSQL’s optimizer to achieve nearly optimal performance, clearly outperforming its built-in cardinality estimator and other alternatives.

PVLDB Reference Format:

Pengfei Li, Wenqing Wei, Rong Zhu, Bolin Ding, Jingren Zhou, and Hua Lu. ALECE: An Attention-based Learned Cardinality Estimator for SPJ Queries on Dynamic Workloads. PVLDB, 17(2): 197 - 210, 2023. doi:10.14778/3626292.3626302

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/pfl-cs/ALECE>.

1 INTRODUCTION

A cardinality estimator in a DBMS [18, 47] estimates the number of result elements of a SQL query before query execution, and thus helps the query optimizer to generate good query plans. In the past, the mainstream of cardinality estimation has always been statistical data-driven methods. Such methods condense information about data into lightweight summaries, *e.g.*, histograms, sketches and data distribution approximation, and adopt analytic functions with the summaries as the input to estimate cardinalities of SQL

† Corresponding authors.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 17, No. 2 ISSN 2150-8097. doi:10.14778/3626292.3626302

queries [27, 47, 67]. However, real-world datasets are often complex and the analytic functions are usually not powerful enough to build correct mappings between coarse data summaries and SQL query cardinalities. Also, SQL queries often contain join predicates but it is difficult and time-consuming to build particular summaries for each join. Computing joint data distributions is also usually intractable due to high computation and storage overhead.

Recently, traditional cardinality estimators have been disrupted by estimators based on learned models. Data-driven models [26, 55, 63, 64] learn tighter data distributions from the underlying database and use analytic expressions to estimate the cardinalities. In contrast, query-driven models [32, 65] utilize the feedback of executed queries in a supervised fashion. The latter learn the relation between cardinalities and query distributions, without paying particular attention to the underlying database. However, neither kind of models can fully make use of both data and queries. It is difficult for them to extract individualized useful information for different queries. A few models [14, 32, 41] consider both data and queries. However, they either only use simple and trivial data information and requires sampling operations over relations [32, 58], or do not support processing queries with joins [14] or complex joins [41].

In addition, existing models have a more critical problem: They do not perform well on **dynamic workloads that mix queries and data manipulation statements including inserts, deletes and updates**. Such statements tend to make estimations difficult as they influence the data distribution and shift the mapping between true cardinalities and query distributions. When the underlying data changes, the joint data distributions among relations and attributes as well as the mapping between queries and true cardinalities also become different. Thus, pure data- or query-driven methods can hardly work on dynamic workloads. Some methods [32, 58], although they consider both data and queries, will also have degraded performance on dynamic workloads as their required featurization or sampling approaches does not support model training and inference with data updates. More importantly, existing methods do not answer **how to reasonably link SQL queries and the underlying data** and build an appropriate mapping among the true cardinalities, queries and data—especially when data is dynamic.

To address these drawbacks, we design an Attention-based LEarned Cardinality Estimator (ALECE) for select-project-join (SPJ) queries. Fig. 1 depicts ALECE in the context of DBMS’s query execution. ALECE is both data- and query-driven. When estimating an SPJ query’s cardinality, it losslessly featurizes the query into a vector. Meanwhile, it efficiently featurizes the current underlying data in




Figure 1: ALECE-based Query Execution.

the database into a set of vectors, called *DB states*, which ‘compress’ the whole database. Both query and data featurizations are of low space overhead and can be efficiently computed. On top of the DB states and query featurizations, ALECE builds a neural network based model to create reasonable connections between them. The model integrates the information of the DB states and query featurization, and feeds them into a feed-forward regression neural network to make the estimates. Roughly speaking, ALECE first learns to assign different weights to the raw DB states X , with each weight showing the correlation between two elements in X . This correlation is an useful distribution information to build suitable mapping between the cardinality of a SQL query with the underlying data. Then, X is mapped into another set of vectors Z which are the weighted combinations of X and better represent the underlying data. ALECE also learns an another weight for each mapped vector z_i in Z and the query featurization q to measure the influence of z_i on q . The weighted combination of Z is a convolution of the DB states and the query featurization. The combination vector is finally used to generate the cardinality estimate.

Our ALECE’s design encounters two challenges. First, we need to build the DB states suitable for dealing with data changes in dynamic workloads, and meanwhile make them efficient to access. To this end, we propose a simple yet effective data featurization approach that is a good approximation of data distributions, sensitive to data changes, and computation-efficient. This approach constructs succinct summaries of the underlying data, *i.e.*, DB states, based on the histogram of each attribute of the database relations. Each time a record is inserted, deleted or updated, we only need to modify the DB states’ vectors relevant to the changed relation. Also, the basic single attribute distributions and even joint distributions are covered by the DB states. These factors together enable us to process dynamic workloads with the DB states. Besides, depending on the requirement of the distribution approximation precision, the number of bins in a histogram can be flexibly adapted. Moreover, other useful information relevant to underlying data can be seamlessly integrated if needed.

Second, we need to extract the implicit relevance between SQL queries and corresponding DB states, and make the information helpful for cardinality estimation. To this end, we adopt the attention mechanism [6, 30, 53] in our model to draw global dependencies between SQL queries and underlying data. The attention mechanism is widely used in a variety of tasks including question answering [25, 50]. Generally, it simulates the process of selection from a set using an attention function that takes as input two main components: a set of queries and a set of key-value pairs. It figures out in an individualized manner which parts of the data play more important roles for different queries, assigns higher weights to the more important and relevant keys for each query, and outputs the

combination of the weighted values. Unlike those concepts in a database, a query in attentions is a specific element for which we need to learn a representation, the role of keys is to respond more or less to the query, and the values are used to compose an answer. Nevertheless, the selection process exactly matches our settings where the SQL queries and underlying data are analogies of the queries and key-value pairs in attentions, respectively.

There are two modules in our ALECE where the attentions are used in different ways. On the one hand, the ‘data-encoder’ module uses a self-attention whose inputs of queries, keys and values all come from the DB states. The self-attention allows the DB states, which correspond to different attributes, to interact with each other. By using the self-attention, the data-encoder module learns the implicit joint distribution information among attributes and computes a smarter representation of the underlying data. On the other hand, in the ‘query-analyzer’ module, the queries set of the attention is exactly a set covering only one featurization vector of a SQL query, while the keys and values come from the output of the data-encoder module. The query-analyzer module outputs a fixed-dimensional ‘answering’ vector integrating the information from the query and data representations. We then use a simple linear regression model to map the answering vector to a cardinality estimate.

Compared to the state-of-the-art cardinality estimation methods, ALECE is able to make more reasonable use of both queries and underlying data. With the help of the two attentions, it answers the questions that ‘which parts of data should a SQL query pays more attention to?’ and ‘how to find the more important data?’ A SQL query usually focuses on some local parts of selected attributes. Also, the join conditions make particular tuples contribute more to the cardinality. Moreover, ALECE is able to adapt to dynamic workloads. In practice, learned models need to be trained with past queries and corresponding DB states, and estimate cardinalities for future queries. The performance of existing query-driven models often dramatically degrades when making predictions on a dynamic database. In contrast, ALECE can make immediate and suitable reactions to data changes by modifying the DB states, and learn an appropriate but implicit mapping between the true cardinality and the query featurization accompanied with the corresponding DB states. Our experimental results show that ALECE is able to make accurate estimates even when the distribution of the underlying data changes. Thus, ALECE is less sensitive to data changes.

In our evaluation, ALECE achieves the best cardinality estimation performance on multiple dynamic workloads. Experimental results show ALECE improves the average end-to-end query time by up to 2.7× faster on the benchmark workload, very close to the optimal results acquired by using true cardinalities. This demonstrates that our ALECE makes more accurate cardinality estimates and helps the query optimizer find better query plans.

We make the following major contributions in this paper:

- We propose necessary principles for a method to featurize the underlying database data and SPJ queries. Accordingly, we design a featurization schema to losslessly featurize an SPJ query and make a reasonable compression of the data. The featurizations can be efficiently updated to support dynamic workloads.
- Based on the featurizations of queries and data, we propose an attention based learned cardinality estimator ALECE, together with detailed analyses.

- ALECE is designed to be a ‘whitebox’ which not only gives estimates but also clear rationale to integrate the SPJ queries and underlying data together in processing dynamic workloads.
- We experimentally validate ALECE’s advantages over more than half dozen representative alternatives on real datasets.

The rest of the paper is organized as follows. Section 2 gives the preliminaries. Section 3 presents the featurizations of data and queries. Section 4 elaborates on ALECE, followed by an analysis of it in Section 5. Section 6 reports on the experimental studies. Section 7 reviews the related work. Section 8 concludes the paper. In addition, due to space limit, we introduce our developed benchmark, which integrates ALECE into PostgreSQL’s query optimizer, and more experimental analyses in an extended version [36].

2 PRELIMINARIES AND PROBLEM

Table 1 lists important notations used in the paper.

Table 1: Notations

R_i	A relation in the database
A_j^i	The j th attribute of the relation R_i
N, T	The number of relations/attributes in the database
$X = \{x_i\}_{i=1}^T$	The set of data featurizations (a.k.a. DB states)
d_x	The number of histogram bins (dimensionality) for a DB state
$q = \langle q_J, q_F \rangle$	A SQL query and its vectorized featurization
d_q	The dimension of a query featurization vector
n_{enc}, n_{ana}	The number of attention layers in the data-encoder/query-analyzer module
K, V, Q	The input keys/values/queries of an attention function

2.1 Cardinality Estimation Problem

Suppose a database D has a set of relations $\{R_1, \dots, R_N\}$. A relation R_i has n_i attributes, i.e., $R_i = (A_1^i, \dots, A_{n_i}^i)$. Each attribute A_j^i can be either *categorical* or *numerical*: the domain of a categorical attribute is a finite set and can be 1-to-1 mapped to an integer set $\{1, \dots, \max_j^i\}$; the domain of a numerical one is $[\min_j^i, \max_j^i]$.

Problem Formulation. Given a SQL query q and a dynamic database D , we want to estimate the *cardinality of q* , denoted as $c(q, D)$, i.e., the number of resulting tuples when q is executed on D .

In this paper, we focus on select-project-join SQL queries with conjunctive filter predicates; the cardinality $c(q, D)$ is the number of tuples after joins and filters, as the following counting query:

$$c(q, D) : \text{SELECT COUNT(*) FROM } R_{i_1}, \dots, R_{i_n} \quad (1)$$

WHERE join predicates J AND filter predicates F

where q involves n relations R_{i_1}, \dots, R_{i_n} , with a set of join predicates J which is a conjunction of join conditions each in the form of “ $R_i.A_x^i = R_j.A_y^j$ ”, and a conjunction of filter predicates F . This formulation allows us to support not only PK-FK joins but also more general joins by specifying join predicates on pairs of joinable attributes (which may or may not be primary/foreign keys) in J . A filter predicate is an relational expression in the form of “ $R_i.A_j^i \text{ op const}$ ” where $\text{op} \in \{<, \leq, >, \geq, =\}$ and const is a fixed value. In q , an attribute can appear in a join or a filter predicate, or both. The support for LIKE predicates is left for future work.

ALECE in the Optimization of SPJ Queries. The estimation results for counting queries in the format of (1) can be used to support the optimization of more complex queries, e.g., widely-used **SPJ** (select-project-join) queries in the following format:

$$q_{\text{SPJ}} : \text{SELECT AGG}_1, \dots, \text{AGG}_m \text{ FROM } R_{i_1}, \dots, R_{i_n} \quad (2)$$

WHERE join predicates J AND filter predicates F

ORDER BY attribute_set_1 GROUP BY attribute_set_2

where each AGG_k ($k = 1, \dots, m$) is an aggregate function over one or multiple attributes which can be COUNT, AVG, MIN and MAX, etc, or can be simply omitted. The join predicate set J and filter predicate set F carry the same meanings with that in (1).

To search for the best execution plan, the query optimizer of a modern DBMS like PostgreSQL first decomposes q_{SPJ} into a series of sub-queries (implicitly) in some fixed order [22]. The cardinalities of these sub-queries are then estimated with the built-in estimator. Accordingly, candidate query execution plans are enumerated and their estimated execution costs given the cardinality estimates are calculated using also a fixed cost model. The plan with the smallest estimated cost is chosen to execute the query physically. Apparently, the execution performance of a query q_{SPJ} is basically determined by the cardinality estimates of its sub-queries. The design of our ALECE enables it to provide more accurate estimates. Our developed benchmark can plug external cardinality estimator into the optimizer to replace the built-in one. Thus, ALECE is able to improve the cardinality estimation for q_{SPJ} ’s sub-queries, and further enable the query optimizer to select a good execution plan.

ALECE is applicable for optimizing even more complex queries. For the sub-queries that it supports, it gives better cardinality estimates; for the sub-queries that ALECE does not support, the optimizer can still use its default cardinality estimator.

Estimation Model on Dynamic Workloads. In reality, the data in a DBMS is seldom static but often keeps being updated. Thus, it is beneficial to design cardinality estimators able to make accurate estimates for a *dynamic workload*, i.e., a sequence of SQL statements including queries, inserts, deletes and updates. Our ALECE aims to support dynamic workloads and provide up-to-date cardinality estimates for queries at any time during the workload.

For a learned cardinality estimator to work on frequent changes of the underlying data distribution without retraining, a straightforward idea is to use the database D itself as part of the input features to train an estimation model. However, this is infeasible as the size of D can be huge and varies continuously. Instead, we use succinct summaries of the database (e.g., fixed-size histograms), called *DB states*, as part of input features to train ALECE. As the database D is updated, the DB states should be updated accordingly (and efficiently) such that they can be fed into the trained model to produce cardinality estimates. Details about featurizing D as DB states are in Section 3.1. We assume that D ’s schema is static. The support for dynamic schema is left for future work.

2.2 Overview of ALECE

An overview of ALECE’s model structure and its role in the query engine is shown in Fig. 1. Features from dynamic data and queries are decoupled and handled by two modules in ALECE. The *data-encoder* module adopts a self-attention structure on DB states to figure out the correlations among all attributes and to learn their joint distribution, whereas the *query-analyzer* module employs a data-query cross attention to discover correlation between the data-encoder’s outputs and the cardinalities of (sub-)queries. The cardinality estimates eventually produced by the query-analyzer in ALECE depend on both data (DB states) and queries.

Offline Training. Training ALECE needs a dataset of queries, their true cardinalities, and the corresponding database information when these queries are executed. The training dataset is obtained

by collecting the true cardinalities of the historical queries executed on a dynamic database for a period of time. We start from featurizing the initial database by generating a set of vectors with fixed dimensionality, *i.e.*, DB states (Section 3.1). Statements in the SQL workload are sequentially processed. For the insert, delete and update statements, we modify the DB states accordingly. When a query comes, it will also be featurized into a vector with fixed dimensionality (Section 3.2). The query features and the current DB states will be packed together as a training sample with the true cardinality as its label. With sufficient training samples collected, ALECE is trained with gradient descent methods.

Online Estimation. A well-trained ALECE can make online estimates on both static and dynamic workloads. Given a new query, we feed its featurization and the up-to-date DB states into ALECE to get the estimates. If the workload is static, *i.e.*, it contains no data update statements, the DB states are constant. Otherwise, the DB states keep changing and the latest ones will always be used.

3 FEATURIZATIONS OF DATA AND QUERIES

The underlying database data and SQL queries are required to be featurized numerically such that our ALECE can deal with. Any featurization method is able to be flexibly adopted by ALECE as long as it satisfies some principles. First, the underlying data needs to be featurized into a set of fixed-dimensional vectors covering enough distribution information. Second, any SPJ query should be losslessly mapped to a fixed-dimensional vector such that ALECE could better understand it. Also, to effectively support the dynamic workloads, the featurization method is supposed to be efficient and of low storage overhead. Following these principles, we propose an method of featurizing the database data and SQL queries numerically. The details are given in Section 3.1 and 3.2, respectively. Moreover, Section 3.3 discusses the properties of our featurization method and how it helps process dynamic workloads. It is noteworthy that our featurization method is specifically designed for our ALECE and it perfectly aligns with the requirements on inputs to ALECE.

3.1 Data Featurization

In our settings, the data featurization X , also known as the ‘DB states’, is a compression of the whole database, which can roughly describe the data of each attribute and the relationships among them. Our ALECE requires the data featurization to be a set of vectors of the same dimension. Here we use the set of histograms for each attribute as the DB states, *i.e.*, $X = \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$ where \mathbf{x}_i is the histogram of the i th attribute and $T = \sum_{i=1}^N n_i$ is the number of all attributes in the database. How to order the T attributes will be introduced in Section 3.2. This featurization method is simple but powerful and we can efficiently access and update the histograms.

In particular, the values of categorical attributes are first converted to consecutive integers numbered from 1. Given an attribute A_i , we use $\text{dom}(A_i)$ to denote its domain or the converted integer set if A_i is categorical. It is easy to show that $\text{dom}(A_i) \subseteq D(A_i) = [l, u]$ where $l = \inf \text{Dom}(A_i)$ and $u = \sup \text{Dom}(A_i) + \epsilon$ with $\epsilon \rightarrow 0^+$. Then, given a time stamp t and the database data at t , we create a d_x -bin-histogram for each A_i . In particular, let $a = \frac{u-l}{d_x}$ and β_j be the number of A_i 's values in $[l + (j-1) \cdot a, l + j \cdot a)$ for $1 \leq j \leq d_x$, the histogram \mathbf{x}_i for attribute A_i could be easily accessed with $\mathbf{x}_i = [\beta_1, \dots, \beta_{d_x}]$. Our DB states at any time is simply a set of T

elements each being a d_x -dim histogram vector. In practice, each β_i in all histogram vectors will be scaled to the range $[0, 1]$ through a suitable affine transformation. The value of d_x can be flexibly modified according to the complexity of the data distribution. Apparently, a larger d_x will make the data featurization capture more distribution information among the attributes but result in extra time and storage overhead. Usually, when strong correlations exist among attributes, a larger value tends to be used for d_x .

3.2 Query Featurization

Following and extending the existing work [56, 65], we featurize a SQL query q into a fixed length vector \mathbf{q}^1 . It is a simple concatenation of two separately generated parts \mathbf{q}_J and \mathbf{q}_F , which featurize the join predicates J and filter predicates F , respectively.

Join featurization. For the N relations numbered from 1 to N , we use $m_1 = \lceil \log_2(N+1) \rceil$ bits to featurize the id of each relation. Similarly, $m_2 = \lceil \log_2(n_{max}+1) \rceil$ bits can featurize the ids of all attributes in any relation, where $n_{max} = \max(\{n_1, \dots, n_N\})$ is the maximum number of attributes in a relation. Thus, any attribute $R_i.A_j^i$ can be uniquely identified with a binary vector of dimension $m = m_1 + m_2$. The first m_1 -dimensional and the last m_2 -dimensional sub-vectors identify the relation and the attribute, respectively. Then a join predicate P is featurized by a $2m$ -dim binary vector $E_J(P)$ with the first and second half sub-vectors refer to the left and right hand side of P , respectively.

Suppose there are Δ possible join patterns, the join featurization \mathbf{q}_J of a SQL query q is a $(2m \cdot \Delta)$ -dim binary vector, containing Δ $2m$ -dim sub-vectors, indicating which join patterns q covers and featurizing their referred attributes. If the join predicates of q contain the i th join pattern P_i , the i th sub-vector of \mathbf{q}_J equals $E_J(P_i)$. Otherwise, this sub-vector is set to zero. Usually, the value of Δ is not large, nearly linearly related to the number of attributes [59, 67]. Unlike existing work [65] that simply featurizes whether each join condition appears in the query, our featurization way is more compact and incorporates more helpful information about the joins.

In practice, we define the attributes in the left and right hand side of a join predicate are ‘equivalent’ and find all equivalence classes in J . Afterwards, we re-organize the join predicates based on the equivalence classes. Given an equivalence class $C = [A_1, \dots, A_t]$ containing t attributes sorted by their m -dim featurizations, we re-create $(t-1)$ join predicates with the i th one to be $A_i = A_{i+1}$. By doing this to each equivalence class and packing the corresponding join predicates together, we generate a new join predicate set J' . The join featurization is actually performed with J' instead of J . In this way, two equivalent join predicate sets in explicitly different forms will be featurized to be the same. For example, in the following formula, join predicate set J_1 and J_2 both have two same equivalence classes. They will be converted to another set J' .

$$\begin{aligned} J_1 &: (A_1^1 = A_1^2 \text{ and } A_1^2 = A_1^3) \text{ and } (A_2^3 = A_2^2) \\ J_2 &: (A_1^1 = A_1^2 \text{ and } A_1^3 = A_1^2) \text{ and } (A_2^2 = A_2^3) \\ J' &: \underbrace{(A_1^1 = A_1^2 \text{ and } A_1^2 = A_1^3)}_{\text{Equi-class-1}} \text{ and } \underbrace{(A_2^2 = A_2^3)}_{\text{Equi-class-2}} \end{aligned}$$

Filter featurization. We sort all attributes according to their m -dim featurizations and use A_k to denote the k th one among T

¹Without ambiguity, \mathbf{q} denotes both a SQL query and its vectorized featurization.

attributes. Without loss of generality, we assume $D(A_k)$ is $[0, 1)$ for each attribute A_k . Thus, the product space $D(A_1) \times \dots \times D(A_k) \times \dots \times D(A_T) = [0, 1)^T$. Apparently, the filter predicates F are equivalent to a hyper-rectangle $[l_1, u_1) \times \dots \times [l_T, u_T)$ which is a subset of $[0, 1)^T$. In particular, for any filter condition on the attribute A_k , we convert it into an equivalent one in the form like $\sigma_{lb \leq A_k < ub}$. Then, the values of l_k and u_k are set to lb and ub , respectively. Specifically,

$$\begin{aligned} (lb \leq A_k) &\sim (lb \leq A_k < 1), & (lb < A_k) &\sim (lb - \epsilon \leq A_k < 1), \\ (A_k < ub) &\sim (0 \leq A_k < ub), & (A_k \leq ub) &\sim (0 \leq A_k < ub + \epsilon) \\ (A_k = x) &\sim (x \leq A_k < x + \epsilon), & &\text{where } \epsilon \rightarrow 0^+. \end{aligned}$$

Above, \sim denotes the equivalence operator.

Accordingly, the featurization of filter predicates \mathbf{q}_F is a $2T$ -dim vector composed of the boundary points of the search hyper-rectangle, i.e., $E_f = [l_1, u_1, l_2, u_2, \dots, l_T, u_T]$. In practice, each l_i and u_i will be normalized to $[0, 1]$.

Concatenation. By concatenating \mathbf{q}_J and \mathbf{q}_F , we get the d_q -dim featurization vector of SQL query \mathbf{q} . Fig. 2 shows an example.




Figure 2: An example of query featurization.

3.3 Discussions

3.3.1 Featurization properties. As we claimed, our way of featurizing database data and SQL queries has the following properties:

1) **Efficiency.** Building the data featurizations of a static database requires looking over each relation once only. An insert/delete/update statement only influences one relation and modifying the relevant histograms takes $O(v)$ time where v is the number of the records involved. The time complexity of featurizing a SQL query is $O(|J| + |F|)$, which is small and can often be ignored.

2) **Low space overhead.** The DB states only contain $T \cdot d_x$ float numbers. In practice, we usually set d_x smaller than 100. The dimensionality of a query featurization is $2m \cdot \Delta + 2T$. Usually, the value of m is small. In our experiments, m is smaller than 10. Our way of re-generating join predicates based on equivalence classes will usually reduce Δ , the number of possible joins. Also, the joins tends to be performed on attributes with primary keys. According to our observations, Δ is usually less than N^2 and featurizing a query on a 8-relation database requires less than 1000 float numbers.

3) **Stability.** The dimensions of data and query featurizations are fixed, no matter how the database or queries change. This ensures the featurizations can be easily processed with learned models.

3.3.2 Why our featurizations work on dynamic workloads? It is noteworthy that no matter how the underlying data changes, the featurization of a given query will not change. In contrast, the DB states will vary with the change of the data in the database. It is a compression of the whole database and able to catch the distribution characteristics of each attribute. Our model takes both featurizations as input and is able to ‘convolute’ the query featurizations with the DB states. Thus, when the data changes, it can react

properly and give different predictions for the same query with different DB states. The experiments in Section 6 and Section C in the extended version [36] show that our model outperforms other state-of-the-art methods on both static and dynamic workloads.

4 DESIGN OF ALECE

Given the DB states X and a SQL query featurization \mathbf{q} , our ALECE can reasonably discover the implicit relations between them that are required for cardinality estimation. The mystery behind lies in the attention mechanisms [53] twice used in our ALECE. In this section, we will review the motivation of ALECE’s design, introduce the details of ALECE including its two modules processing X and \mathbf{q} , respectively, and the training process of ALECE.

4.1 Motivations and ALECE Overview

An attention function maps a query and a set of key-value pairs to an output, where the query, keys, values, and output are all vectors [53]. Here, the concepts of key-value pairs and queries can be analogous to retrieval systems. Take a user’s search behavior on an e-commerce platform like Amazon as an analogy. When the search engine receives a query (the text in the search bar), it maps the query against a set of keys (item names, tags and descriptions, etc.) associated with values (candidate items) in the database and outputs the best matched items. The output is a weighted sum of the values, where each weight is computed by a compatibility function that measures the relevance between the query and keys.

The idea behind the attention mechanism is to encode the input key-value pairs set, and utilize the most relevant parts of the keys, associated with values, with the query in a flexible manner. Through a weighted combination of all encoded input vectors, this mechanism ‘answers’ the query with the most relevant vectors getting the highest weights. This idea perfectly fits in our research problem as a SQL query’s featurization \mathbf{q} is a natural query vector. Besides, the DB states X can be seen as the item information in the above example and be used as the keys and values in the attention functions. Thus, we extend this idea and design ALECE, which takes full advantage of the attention mechanism to accurately estimate cardinalities of SQL queries on dynamic workloads. Fig. 3 illustrates the structure of ALECE.

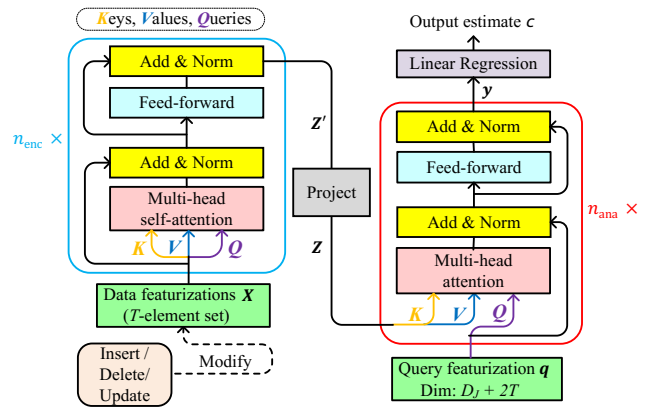


Figure 3: ALECE structure.

Our attention-based model ALECE is designed like a soft lookup table. It fetches information from a better representative of the DB

states X using features from the transformation of q as indices. It is composed of two modules. The left **data-encoder** module maps $X = \{x_1, \dots, x_T\}$ to another set of representations $Z = \{z_1, \dots, z_T\}$. It learns the implicit joint distribution information among all attributes by computing the relevance between any pairs of DB states through a stack of self-attention layers. This distribution information is embodied in the set Z . Given Z , the right **query-analyzer** module adopts another stack of attention layers to measure the relevance between the query featurization q and Z , and generates an ‘answer’ vector y . Finally, vector y is fed into a linear regression layer and q ’s cardinality given DB states X is estimated and returned.

4.2 Attention in ALECE

Attention background. Before showing the data-encoder and query-analyzer modules in ALECE, it is worth briefly introducing the attention mechanism. In neural networks, attention is a technique that is meant to mimic cognitive attention. Its motivation behind is that the network should devote more focus to the important parts of the data instead of treating all data equally. It uses an attention function to discover which parts of the data should be emphasized. The function maps a query and a set of key-value pairs to an output, which is a weighted sum of the values. Usually, the function computes the similarity (relevance) between each pair of query and key with some metric, and uses it to produce the weight assigned to the corresponding value.

Our ALECE uses the ‘Scaled Dot-Product Attention’ [53], namely Attn, as the attention functions in the attention layers from both the data-encoder and query-analyzer modules. The input keys and values are packed together into matrices K and V , of dimension $S \times d_k$ and $S \times d_v$, respectively, where S denotes the size of the original key-value pair set.² The d_k -dim query vector is converted to a $1 \times d_k$ matrix Q . The function Attn computes the dot products of the given query Q with all keys K , divides each by $\sqrt{d_k}$, and applies a softmax function to obtain the weights on the values V :

$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{Q \cdot K^T}{\sqrt{d_k}}\right)V$$

The dot product above is used as the similarity metric. To prevent the dot products from growing too large such that the gradients of the softmax function ‘vanishing’, we divide the dot product by the factor $\sqrt{d_k}$. In practice, the function Attn is computed with batches of queries simultaneously to improve the efficiency. Before feeding the input to Attn, we adopt a multi-head projection mechanism [53] to first project the queries, keys and values multiple times with h different linear projections. This operation is able to enhance the representation ability of ALECE. Due to space limit, more details are given in Section A in the extended version [36].

Attention in ALECE’s two modules. The multi-head attention layer appears twice in ALECE. The first is a self-attention layer in the data-encoder module. It is fed with the same query, key and value matrix X of dimension $T \times d_x$, and outputs a matrix $Z' \in \mathbb{R}^{T \times d_x}$ that will be later projected to another matrix $Z \in \mathbb{R}^{T \times d_q}$. Then, Z is used as the key and value matrix of the other attention layer in the query-analyzer module, where the input query set contains only one SQL query featurization vector. It is noteworthy that

²In the rest of the paper, we regard a set of vectors as a matrix. In other words, we do not distinguish a set of key/value/query vectors with a key/value/query matrix.

the “queries” set in the self-attention layers does not come from the SQL query. In addition, we do not need to order the vectors in the DB states. Thus, unlike transformer-like models to process sequences [7, 13], ALECE does not need the positional encodings [53] or positional embeddings [16] of the attention layers’ inputs.

Table 2 states the input sources and output of the two modules’ attention layers, and what each module is able to learn. The rationale behind them is given in the following two subsections.

4.3 Data-encoder

Suppose a random variable V ’s value is taken randomly with an equal possibility from the possible values of an attribute. As a result, the corresponding histogram can be regarded as the rough distribution function of V . Thus, the DB states consists of a series of ‘marginal’ distributions describing single attributes. However, SQL queries tend to cover join predicates that need to know the joint distribution information among multiple attributes. Usually, the distributions of different attributes are not independent. It is difficult and impractical to directly derive the joint distribution from the marginal distributions. To address this issue, we design the data-encoder module that makes use of attentions, to establish a bridge between the marginal distributions and the joint distribution.

The data-encoder takes as input the DB states X and feeds it into a n_{enc} -stacked layers. Each layer is identical and has two sub-layers. The first is the multi-head self-attention sub-layer N_{sa} which takes the same inputs of keys, values and queries—three matrices equal to X or the output of the last layer, and outputs another T -element set \tilde{Z}_i . On top of N_{sa} , the feed-forward sub-layer FF uses stacked fully connected networks and nonlinear activation functions, e.g., ReLU, to map \tilde{Z}_i into the data representation set Z'_i . Besides, to address the degradation problem and ease the model training, we employ a residual connection [24] around each sub-layer, followed by layer normalization [5], following the settings in [53]. Thus, the output of each sub-layer is $\text{LayerNorm}(V + \text{sub-layer}(V))$. The sub-layer is either N_{sa} or FF. The analytic expression of the output representations Z' is described as follows.

$$Z' = Z'_{n_{\text{enc}}}, \text{ where } Z'_0 = X, \text{ and } \forall 1 \leq i \leq n_{\text{enc}},$$

$$\tilde{Z}_i = \text{LayerNorm}(Z'_{i-1} + N_{\text{sa}}(\overbrace{Z'_{i-1}}^{\text{keys}}, \overbrace{Z'_{i-1}}^{\text{values}}, \overbrace{Z'_{i-1}}^{\text{queries}})),$$

$$Z'_i = \text{LayerNorm}(\tilde{Z}_i + \text{FF}(\tilde{Z}_i)), \text{ with } \text{FF}(V) = \text{ReLU}(WV + b)$$

Finally, the output matrix Z' will be linearly projected to the representation matrix $Z \in \mathbb{R}^{T \times d_q}$. This projection is to align the dimensionality of Z with that of the query featurization q .

The keys, values and queries in the attention layers are the same set. They are either the DB states X or the output of the previous stacked layer. This setting makes each element in the output set of a layer attend to all outputs in the previous layer and thus attend to all DB states. More importantly, the self-attention layer quantitatively ‘calculates’ the relevance between a pair of elements from any two histograms. It is noteworthy that each element of a histogram describes the local distribution of an attribute. Therefore, the data-encoder module is able to more effectively discover the implicit connections between any pair of DB states, and thus exhibits behavior in relation to the joint distribution of multiple attributes in the output set. Compared to other neural network architectures like multilayer perceptron (MLP) [23], self-attention

Table 2: The overall picture of ALECE’s two modules, given DB states X and SQL query featurization q

Data-encoder (Self-attention)			Query-analyzer (Data-query cross attention)		
Input source	What to learn?	Output	Input source	What to learn?	Output
Keys: X	Relevance among the DB states X ,	Another vector set Z cov-	Keys: Z	Relevance between q with the ele-	Final ‘answering’ vector y wh-
Values: X	and thus the joint distribution inf-	ering joint distribution	Values: Z	ments in Z , showing which parts	ich will be directly turned into
Queries: X	ormation of multiple attributes.	information of attributes	Queries: q	of data are more important.	the cardinality estimate c .

could yield more interpretability and have higher representative abilities. Through self-attention layers, we create links among all DB states, or equivalently, all attributes in the database. After fine-tuning the parameters of the self-attention and feed-forward layers, the relationship information helpful to the cardinality estimation task is implicitly covered and encoded into the output Z of the data-encoder module. This information will be processed and utilized in the query-analyzer module.

4.4 Query-analyzer

The query-analyzer module attempts to discover and measure the relevance, through data-query cross attention layers, between the SQL query and each element of Z , the output of the data-encoder module covering joint distribution information among attributes.

This module is also a stacked structure composed of n_{ana} identical layers. Similar to the data-encoder module, each layer here is composed of a multi-head attention sub-layer and a fully connected sub-layer. Also, residual connections are employed around each sub-layer, followed by layer normalization. Unlike the data-encoder module, the input sets of key-value pairs and queries to the ‘data-query’ attention sub-layer here are not from the same place. We use Z , the output of the data-encoder module, as both the input key and value matrices, while the query set here comes from either the query featurization or the output of the previous layer. It is noteworthy that the input query set and the output set of each attention sub-layer have only one d_q -dim element.

The data-query attention sub-layer establishes a bridge between the queries and the data. It individualize each SQL query featurization and presents different ‘answers’ by enhancing the influences of some parts of the input key-value pair set Z , while diminishing other parts. Learning which part of the data is more important depends on the relations between queries and keys, and this is measured with the attention functions. Suppose a SQL query contains join predicate $R_1.A_1^1 = R_2.A_2^2$ and filter predicate $R_1.A_2^1 > 1$, and the attributes $R_1.A_1^1$, $R_2.A_2^2$ and $R_1.A_2^1$ are numbered i , j and k , respectively. The data-query attention sub-layer will pay more attention to the part of the vectors in the set Z that are relevant to x_i , x_j and x_k . The effect of particular attention can be realized through suitable parameters of different layers in both modules.

After accessing y , the output of the final query-analyzer layer, we use a simple linear regression layer LR to calculate a scalar value as the cardinality estimate c . The process of accessing c with the input of Z and q is described as follows.

$$c = LR(y), \text{ where } y = y_{n_{ana}}, y_0 = q, \text{ and } \forall 1 \leq i \leq n_{ana}, \quad (3)$$

$$y'_i = \text{LayerNorm}(y_{i-1} + N_{dq}(\underbrace{Z}_{\text{keys}}, \underbrace{Z}_{\text{values}}, \underbrace{y_{i-1}}_{\text{queries}})),$$

$$y_i = \text{LayerNorm}(y'_i + FF(y'_i))$$

4.5 Training of ALECE

Fine-tuning the parameters of ALECE requires a training dataset of which each element is a 3-tuple (q_i, X_i, c'_i) , where q_i is the query featurization of a SQL query q_i and X_i is the associated DB states

of the dynamic database. By executing q_i on the database of which the DB states is X_i , we will get the true cardinality c'_i , which will be used as the label. In practice, we will take the logarithm of c'_i to make the range of the labels not too large. Collecting the training dataset is not difficult. Usually, we only need to collect the feedback of executed queries on a dynamic database. Then, we will get the training dataset with three lists $\mathcal{X} = [X_i]$, $\mathcal{Q} = [q_i]$ and $\mathcal{C}' = [c'_i]$. They will be split into batches to train our ALECE.

We use the mean-weighted-squared-error function MWSE taking input of the batch card predictions c_b and the true cards c'_b as well as their weights w_b with batch size B as the loss function, i.e., $MWSE(c_b, c'_b) = \frac{1}{B} \sum_{i=1}^B w_i (c_i - c'_i)^2$. The parameters of the linear regression layer, the attention and feed-forward sub-layers in both modules are trained by gradient descent with batches in an end-to-end fashion. Here, the value w_i is proportional to that of $\log c_i$. In particular, $w_i = \frac{\log c_i}{\sum_j \log c_j}$. We use the weight w_i in the loss function because it is usually beneficial to emphasize the queries with larger true cardinalities as their execution times tend to be longer. The procedure of training ALECE using the these lists \mathcal{X} , \mathcal{Q} and \mathcal{C}' is detailed in Algorithm 1.

Algorithm 1: TRAIN-ALECE

Input : Three data lists: $\mathcal{X} = [X_i]$, $\mathcal{Q} = [q_i]$, $\mathcal{C}' = [c'_i]$
The maximum number of training epochs M_e

Output: Well-trained ALECE

- 1 Compute the weight vector $w = [w_i]$ according to \mathcal{C}' .
- 2 **for** $i \leftarrow 0$ **to** M_e **do**
- 3 Train-Epoch($\mathcal{X}, \mathcal{Q}, \mathcal{C}', w$)
- 4 **if** ALECE performs better on validation set **then**
- 5 Assign current parameters to ALECE.
- 6 **return** ALECE whose parameters are fine-tuned.
- 7 **Function** Train-Epoch($\mathcal{X}, \mathcal{Q}, \mathcal{C}', w$):
- 8 Shuffle and split $\mathcal{X}, \mathcal{Q}, \mathcal{C}, w$, generate n_B equal-size batches
[$(\mathcal{X}_b, \mathcal{Q}_b, \mathcal{C}'_b), w_b$], $b = 1, \dots, n_B$.
- 9 **for** $b \leftarrow 0$ **to** n_B **do**
- 10 Compute the vector c_b with Eq (3).
- 11 Compute the MWSE loss given c_b, \mathcal{C}'_b and w_b , and apply gradient descent to tune model parameters.

In practice, a training dataset is split into two parts: the first part is the three data lists used to fit the parameters in the learning process (lines 7-10); the second part is used as the validation set to choose best parameters for ALECE and avoid overfitting (lines 4-5). Also, we will apply the early stopping strategy [44], to stop the training process when the error on the validation set grows.

5 ANALYSIS OF ALECE

In this section, we discuss why ALECE is suitable for dynamic workloads and further analyze its properties.

5.1 ALECE on Dynamic Workloads

It is critical for databases to process dynamic workloads, i.e., mixture of queries and data manipulation operations. This requires a DBMS’ query optimizer to be able to accurately predict the cardinalities of the (sub-)queries performed on a continuously changing database. To reduce the related overhead, the DBMS regularly maintains its

cardinality estimators instead of modifying it immediately after each update. Between two model-maintaining time points, the cardinality estimator model is usually fixed but still expected to make accurate estimates.

The traditional histogram-based cardinality estimators [12, 21, 43, 47] currently used in systems like PostgreSQL [20] make simple and often unreasonable assumptions, like the mutual independence of attributes. Their estimates are inaccurate, especially when the data is dynamic. This renders it difficult for the optimizer to choose good query plans. Existing learned cardinality estimation methods either do not support dynamic workloads [65] or need to completely re-build the model at model-maintaining time points [26, 63, 67]. This process is often too time-consuming to be feasible.

Compared to existing methods [26, 63, 67], the training of ALECE is efficient. Also, it is easily maintained: it can be incrementally updated at model-maintaining time points. Moreover, its architecture design is reasonable and suitable for processing dynamic workloads. It is noteworthy that the SQL queries and underlying data are decoupled in ALECE such that ALECE can learn something whenever the DB states X or the query featurization q in the training dataset changes. Thus, ALECE is able to avoid overfitting certain datasets. Instead, it applies the attention mechanisms on its two modules to learn the properties of the data distribution which ‘generates’ the underlying data, and how these properties influence the cardinality of a SQL query. It is able to reasonably approximate $c(q, DB)$, the true but implicit mapping among the queries, dynamic data and cardinalities. When either the query featurization or DB states, or both of them, get changed, ALECE can still output accurate results. Our experimental results indicate that ALECE achieves good performance even if there is distribution discrepancy between the training and testing data.

A straightforward baseline without attentions. It is mentioned in Section 3.3 that the DB states will be dynamically modified when data is inserted to/deleted from the database. This renders it possible to use a stable well-trained model to predict the cardinalities of queries on a dynamic workload. As the dimensions of DB states and query featurizations are fixed, we can also adopt a straightforward method without attentions to process them. For example, for each pair of data featurizations X_i and query featurization q_i , we flatten the matrix X_i into a vector and concatenate the vector with q_i to generate another vector v_i . Subsequently, v_i and the associated layer are fed into a common supervised neural network like multilayer perceptron (MLP) [23]. However, a straightforward neural network like MLP is usually not powerful enough to discover the implicit relations between SQL queries and the underlying data. Its performance heavily relies on the similarity between the distributions of training and testing datasets. When the workload is static, *i.e.*, the database data never changes, straightforward methods perform well. However, the cardinality estimator model needs to use the current training data it observes to make predictions for ‘future’ data. The distribution of future DB states may be highly different from that of the current data. In contrast, the application of the attention mechanisms in our ALECE make it possible to make accurate estimates even when the distributions of the underlying data change. Experimental results in Section 6 show the great advantages of ALECE over MLP on processing dynamic workloads.

5.2 Overhead of ALECE and Its Extension

Several other good properties make ALECE more practical and attractive as a cardinality estimation method in modern DBMS’s query optimizers. First, the storage and training overhead of ALECE is affordable. On the one hand, the sizes of ALECE on the three datasets in our experiments are both smaller than 23 MB. On the other hand, training ALECE from scratch requires less than 12 minutes. Also, the estimation latency of ALECE is less than 11 ms. These overheads make ALECE able to process real world workloads.

Second, different from FLAT [67] and DeepDB [26], ALECE directly estimates the cardinality instead of the selectivity of a SQL query. To get the cardinality, the selectivity estimation methods usually needs to estimate the size of the corresponding join table with sampling-based techniques [66]. However, when multiple relations are involved, the statistical variance often becomes large and results in highly inaccurate estimates.

Last but more importantly, ALECE establishes a general framework that is not limited to cardinality estimation tasks. By replacing $COUNT(*)$ with other aggregation functions, the special aggregate analytic queries of Format (2) can be further transformed to more general ones. As a matter of fact, ALECE can be easily extended to approximate the results of more general aggregate analytic queries. When processing queries with other kinds of aggregate functions, we only need to make slight modifications to the data and query featurizations, including featurizing the extra aggregate functions and optional GROUP BY clauses in the query featurizations, and incorporating more data description information specific to the aggregate queries. We left the support of the general aggregate analytic queries for future work.

6 EXPERIMENTAL STUDIES

This section reports the experiments that compare ALECE with selected alternatives. All methods are implemented in Python 3.9 and evaluated on a Linux server with a 96-core Xeon(R) Platinum 8163 CPU and 376GB memory. The implementation of ALECE is open [1]. Due to space limit, we present additional experimental results and analyses in Section C in the extended version [36].

6.1 Experimental Settings

Datasets. We use three real datasets to evaluate all models.

- **STATS** contains 8 relations (*users, posts, postLinks, postHistory, comments, votes, badges, tags*) with 43 attributes [2]. There are 1,029,842 records in total. An existing open workload with 146 queries are marked as ‘testing queries’. They are associated with 2,603 sub-queries. We created another 1000 different queries with sub-queries, which are used as the training and validation data.
- **Job-light** is a subset of the IMDB dataset [3]. It contains 6 relations (*cast_info, movie_info, movie_companies, movie_keyword, movie_info_idx, title*) with 14 attributes. There are 62,118,470 records in total. The testing query set contains 208 queries, associated with 3,248 sub-queries. Similarly, another 2,000 queries as well as their sub-queries are created for training the models.
- **TPCH** (1 GB) [4] is a widely-used benchmark dataset of a suite of business oriented relations (*customer, lineitem, nation, orders, part, partsupp, region and supplier*). We remove the string type ‘comment’ attribute from each relation as it is not supported by ALECE or other methods. We use the remaining 46 attributes.

There are 8,661,245 records in total. We randomly create 123 testing queries and 1,554 testing sub-queries for evaluation. Another 2,000 queries and their sub-queries are used for training.

The join information among relations in these datasets are shown in Fig. 5 in the extended version [36]. All joins in the queries on the Job-light dataset are PK-FK joins, while the queries on the other datasets involve more complex many-to-many joins.

Dynamic workloads. For each dataset, we create three different dynamic workloads, each of which is a random mix of inserts, deletes, updates and query statements. These workloads are differentiated according to the ratios among the numbers of inserts, deletes and updates, and the distributions of the inserted records:

- **Insert-heavy:** #inserts : #deletes : #updates = 2 : 1 : 1.
- **Update-heavy:** #inserts : #deletes : #updates = 1 : 1 : 2.
- **Dist-shift** is a special Insert-heavy workload but having skewed distribution of the inserted records. In particular, the inserted records are selected intentionally such that their first attributes' values are the first 30% smallest among all data.

To generate a dynamic workload, we randomly select about $\frac{2}{3}$ of the records as the initial datasets to bulk load all the relations. Subsequently, the insert and update statements in a dynamic workload will insert the remaining $\frac{1}{3}$ of the records to the database, while the delete and update statements will remove or change some records. For simplicity and clarity, we assume that each delete or update statement only influences one record. These statements are equally split into two parts: the former 'training' part and the latter 'evaluation' part. To ensure that ALECE learns from enough features, we make three copies of the training queries and their sub-queries and randomly mix them with the training part. The training part of the workload is used as the training dataset for building all cardinality estimation methods. Note that the true cardinalities of the queries in the training part are also available to all methods.

The way of mixing testing queries with the data manipulation statements in the evaluation part are slightly different. For each testing query, we pack it up with its sub-queries together. Then, those packs are shuffled and randomly mixed with data manipulation statements in the evaluation part. Also, we'd like to know how cardinality estimators perform when a certain amount of data in the database changes. Thus, it is assumed that when each testing query is executed, the associate **changing rate** ρ of the underlying database data is larger than a pre-defined threshold. We use H to denote the set of all records in the database after executing all the statements in the training part. Given any query q_i in the evaluation part, suppose B_i is the set of all records in the database when q_i is executed, and the numbers of inserts, deletes and updates from the evaluation part so far are I_i , D_i and U_i , respectively. The changing rate ρ_i is defined as $\rho(t) = \frac{|H \Delta B_i|}{|H|} = \frac{I_i + D_i + 2U_i}{|H|}$, where Δ denotes the symmetric difference operation of two sets. When there is no ambiguity, the subscript i is omitted.

Both training and testing queries are randomly generated. In particular, to generate a query, we run a series of Bernoulli tests with $p=0.5$ to determine which relations appear. Then, we enumerate all possible subsets of the join conditions among the selected relations such that they can be connected by these join conditions. One of the subsets is randomly selected as the join predicate set J . The attributes in the filter predicates F and their operators (\leq , \geq , etc)

are determined with the analogous Bernoulli tests. The right hand side values of the filter predicates are randomly sampled from the initial dataset.

Competitors. We include the following representative methods:

- **PG** is the simplest 1D histogram based cardinality estimation method used in PostgreSQL [20].
- **Uni-Samp** [37] uniformly samples a set of tuples with a given probability p for each relation. The cardinality estimate of a query q equals N_q/p^m where N_q is the number of returned tuples by executing q on the sample sets and m is the number of involved relations in q . We tune the value of q to make Uni-Samp's storage overhead or latency is comparable to that of the other methods. The values of the sampling ratio p for the STATS, Job-light and TPC-H datasets are set to 0.1, 0.06 and 0.04, respectively.
- **DeepDB** [26] is based on a Sum-Product-Network (SPN) [42]. It learns a pure data-driven model to capture the data's joint probability distribution. Following the same settings in [26], we set the RDC independence threshold to 0.3 and split each SPN node with at least 1% of the input data.
- **FLAT** [67] improves SPN based on factorize-sum-split-product network (FSPN) [62], a graphical model, to adaptively model the joint probability density function of attributes. It is data-driven.
- **NeuroCard** [63] is a data-driven method, extending Naru [64] into the multi-table case, while Naru is a Deep Auto-Regression (DAR) [19] based cardinality estimation algorithm for a single table. The sampling size of the NeuroCard model is set to 8,000, following the settings in the paper.
- **FactorJoin** [59] combines classical join-histogram methods with learned single table cardinality estimates into a factor graph.
- **MLP** [46] is the baseline neural network based method introduced in Section 5.1.
- **MSCN** [32] is a multi-set convolutional network which adopts the information from both queries and data.
- **NNGP** [65] adopts the Neural Network Gaussian Process (NNGP) model [34] to learn from queries as well as their true cardinalities in a supervised manner.

In addition, we also include the comparison with the results generated by true cardinalities. This **Optimal** method measures the best performance a method can achieve.

These competitors are chosen because they have better overall performance over other statistical and learned cardinality estimators. The comparisons are reported in benchmark and evaluation papers [22, 51] and other existing works like [60, 63, 67]. Table 3 summarizes the properties of all methods. The better performance is indicated in bold.

Table 3: Properties of different methods

Method	Data-driven	Query-driven	Building time	Space overhead	Suitable for dynamic workload
PG [20]	✓	×	small	low	×
Uni-Samp [37]	✓	×	small	medium	×
DeepDB [26]	✓	×	medium	medium	×
FLAT [67]	✓	×	large	high	×
NeuroCard [63]	✓	×	medium	medium	×
FactorJoin [59]	✓	×	medium	small	×
MLP [46]	×	✓	small	low	×
MSCN [32]	✓	✓	small	small	×
NNGP [65]	×	✓	small	medium	×
ALECE (ours)	✓	✓	small	medium	✓

Table 4: Overall performance of cardinality estimation models on dynamic workloads

Data	Model	Insert-heavy											Update-heavy										
		E2E	Q-error				P-error				Size (MB)	Building Time (Min)	Latency (ms)	E2E	Q-error				P-error				
		Time(S)	50%	90%	95%	99%	50%	90%	95%	99%				Time(S)	50%	90%	95%	99%	50%	90%	95%	99%	
STATS	PG	7,790	190	1.4·10 ⁹	1.1·10 ⁶	1.8·10 ⁷	2.60	25.44	41.23	243	-	-	-	4,337	524	43,096	3.7·10 ⁵	1.1·10 ⁷	1.78	19.52	29.01	178	
	Uni-Samp	6,646	1.35	12.47	>10 ¹⁰	>10 ¹⁰	2.71	16.77	21.65	201	3.34	0.02	239	6,002	1.27	9.3·10 ⁹	>10 ¹⁰	>10 ¹⁰	1.10	12.86	33.19	153.11	
	NeuroCard	>30,000	17.37	1,388	7,402	3.0·10 ⁵	2.38	39.18	824	84,415	90.76	23.69	32.42	22,193	18.49	1,252	6,784	3.1·10 ⁵	2.10	31.75	347	11,730	
	FLAT	>30,000	12.77	1,979	12,897	8.6·10 ⁵	2.22	70.08	2,202	8,738	210.33	53.77	13.67	24,319	16.00	2,376	17,811	9.4·10 ⁵	2.20	29.76	724	1.4·10 ⁵	
	FactorJoin	>30,000	22.62	2,593	31,936	1.6·10 ⁶	2.35	66.12	876	8,384	1.64	0.42	1.33	>30,000	23.76	2,939	36,198	2.4·10 ⁶	2.75	115	912	4,409	
	MLP	>30,000	2.4·10 ⁶	>10 ¹⁰	>10 ¹⁰	>10 ¹⁰	7.35	491	2,294	6,887	8.63	3.11	3.17	>30,000	2.3·10 ⁶	>10 ¹⁰	>10 ¹⁰	>10 ¹⁰	6.66	502	3,025	11,584	
	MSCN	27,758	20.09	2,870	17,037	2.6·10 ⁵	2.46	87.23	736	8,738	1.61	11.87	1.02	25,766	18.89	2,382	16,947	1.8·10 ⁶	2.56	64.64	246	46,721	
	NNGP	12,883	9.88	827	4,652	2.6·10 ⁵	1.41	5.31	27.33	3,588	32.44	0.89	32.66	11,985	8.15	567	3,639	2.5·10 ⁵	1.52	5.62	13.81	2,753	
	ALECE	2,901	1.29	5.07	11.54	62.52	1.07	1.34	1.59	2.32	22.31	9.25	8.25	2,402	1.42	5.06	10.40	44.69	1.08	1.45	2.13	2.86	
	Optimal	2,871	1	1	1	1	1	1	1	1	1	-	-	2,349	1	1	1	1	1	1	1	1	1
Job-light	PG	6,128	1.59	7.14	20.85	223	1.10	1.32	1.48	2.89	-	-	-	4,747	1.67	6.20	14.09	70.13	1.10	1.39	1.64	3.06	
	Uni-Samp	6,401	1.23	2.59	6.53	>10 ¹⁰	1.08	1.85	2.79	146	23.95	0.52	163	4,491	1.33	3.61	7.50	>10 ¹⁰	1.12	1.71	2.26	3.67	
	DeepDB	26,636	11.31	663	6,328	5.4·10 ⁵	1.66	13.71	26.01	49.57	11.52	11.93	11.92	23,190	10.64	628	4,883	8.8·10 ⁵	1.50	11.81	27.71	135	
	NeuroCard	27,744	14.73	978	6,766	4.7·10 ⁵	1.81	12.40	37.25	144	42.93	9.11	15.90	16,553	13.49	1,000	6,489	1.5·10 ⁶	1.60	15.10	25.94	50.68	
	FLAT	23,876	11.22	626	5,786	4.4·10 ⁵	1.66	13.98	27.26	89.67	11.10	17.83	3.96	22,531	10.13	490	4,221	1.1·10 ⁶	1.42	12.87	34.54	195	
	FactorJoin	16,665	25.14	3,456	27,938	2.5·10 ⁶	1.52	5.45	10.23	47.09	4.66	26.22	1.33	13,071	27.31	6,134	36,155	2.2·10 ⁶	1.42	4.83	9.43	22.75	
	MLP	5,610	4.30	23.14	59.37	2,411	1.13	1.70	2.10	3.12	6.19	2.68	2.28	5,351	5.59	62.02	138	1,441	1.17	2.02	2.88	4.09	
	MSCN	26,665	24.42	1,809	9,217	1.4·10 ⁵	2.07	20.98	28.58	106	1.56	33.37	0.77	29,118	26.77	1,448	7,980	1.5·10 ⁵	2.59	19.87	61.18	257	
	NNGP	15,430	4.5·10 ⁸	>10 ¹⁰	>10 ¹⁰	>10 ¹⁰	3.57	8.39	11.97	24.47	32.44	0.83	29.86	4,959	4.99	59.86	190	50,036	1.20	2.05	2.99	4.96	
	ALECE	5,168	1.14	2.08	3.85	14.82	1.07	1.24	1.32	1.77	22.25	7.73	7.28	3,839	1.09	1.93	3.63	16.28	1.07	1.23	1.29	1.57	
Optimal	5,063	1	1	1	1	1	1	1	1	1	-	-	3,792	1	1	1	1	1	1	1	1	1	
TPCH	PG	3,230	1.35	4.88	7.90	16.16	1.44	1.57	1.59	1.62	-	-	-	2,385	1.33	5.01	6.87	14.64	1.41	1.56	1.59	1.62	
	Uni-Samp	>30,000	1.22	2.94	>10 ¹⁰	>10 ¹⁰	1.81	2.22	45.60	514	23.29	0.17	24.52	>30,000	1.30	3.41	>10 ¹⁰	>10 ¹⁰	1.75	19.35	50.34	211	
	NeuroCard	10,616	43.43	7,071	27,735	3.5·10 ⁵	3.75	10.91	22.48	106	704.98	34.38	34.77	6,562	45.73	5,896	20,931	1.2·10 ⁵	3.16	11.04	16.49	245	
	FLAT	11,428	46.04	12,236	70,383	8.4·10 ⁷	3.56	10.91	25.71	1,182	168.65	41.42	10.06	7,864	42.32	7,328	26,470	3.7·10 ⁵	3.56	15.82	21.70	425	
	MLP	7,261	51,475	1.4·10 ⁷	3.5·10 ⁷	4.7·10 ⁹	4.71	14.11	22.74	36.09	8.63	3.57	3.62	8,342	79,317	3.7·10 ⁷	9.1·10 ⁷	>10 ¹⁰	5.12	17.63	32.74	63.39	
	MSCN	7,718	36.36	6,709	26,744	4.4·10 ⁵	3.58	15.11	25.71	387	1.58	23.43	0.79	8,353	41.66	5,457	21,433	1.7·10 ⁵	3.80	14.86	25.69	472	
	NNGP	18,618	1.3·10 ⁹	>10 ¹⁰	>10 ¹⁰	>10 ¹⁰	5.56	47.40	59.18	99.55	32.44	0.92	41.76	3,128	25.88	432	986	4,968	1.54	7.98	10.90	20.66	
	ALECE	3,233	1.15	2.29	3.37	8.86	1.43	1.57	1.60	1.62	22.34	11.62	10.35	2,380	1.25	2.11	2.95	10.87	1.41	1.56	1.59	1.62	
	Optimal	3,227	1	1	1	1	1	1	1	1	1	-	-	2,378	1	1	1	1	1	1	1	1	1

Evaluation Metrics. We use three metrics to evaluate all methods:

- **E2E time** is the total execution time of all queries in the evaluation part of a workload, in which we feed the query optimizer the estimated cardinalities of all relevant sub-queries. Those estimation results are acquired using cardinality estimation methods. It is the most important evaluation metric as it directly connects to the query optimizer and objectively shows if a cardinality estimation method could help improve the query performance of a DBMS. This type of evaluation requires an improved benchmark over the existing work [22]. This benchmark can integrate the estimation results by an external method to PostgreSQL. We show its details in Section B in the extended version [36].
- **P-error** [22] measures the gap between the optimal query plan and the generated plan based on the estimated cardinalities, without executing the given query. In particular, given a query q , a query plan P and the set c of estimated/true cardinalities of q 's all sub-queries, the DBMS will output an estimated cost $C(P, c)$ with a cost function C . By feeding the query optimizer c^T , the set of true cardinalities of q 's all sub-queries, we can get the optimal plan P^T for q . Similarly, a cardinality estimation method A outputs for q a set c^E which results in another query plan P^E . Then, $P\text{-error}(A, q) = \frac{C(P^E, c^T)}{C(P^T, c^T)}$, where the denominator $C(P^T, c^T)$ is the optimal execution cost, and the numerator $C(P^E, c^T)$ is the cost by feeding the true cardinalities of sub-queries to the query plan generated by method A. In other words, $C(P^E, c^T)$ is the actual execution cost of q if method A is adopted.
- **Q-error** [39] measures the distance between the estimated cardinality P and the true cardinality T of a query. In particular, $Q\text{-error}(P, T) = \max(\frac{T}{P}, \frac{P}{T})$.
- **Storage overhead** is the memory size used by a method.
- **Building time** indicates the offline training time of query-driven methods or construction time of data-driven methods.
- **Estimation latency** is the average estimation time per sub-query used by a cardinality estimation method.

As the optimizer only requires the cardinality estimates of the sub-queries, our evaluations also involve the testing sub-queries only.

Parameter Settings. We use a 40-bin-histogram for each attribute, *i.e.*, $d_x = 40$. The values of n_{enc} and n_{ana} , *i.e.*, the number of stacked attention layers in the data-encoder and query-analyzer, respectively, are both set to 4. To train our ALECE, we use an Adam optimizer [31] with a learning rate of 0.01 and a batch size of 128.

6.2 Performance on Dynamic Workloads

For each dynamic workload, we build all methods using the data in the training part to make estimates for the testing sub-queries in the evaluation part. The featurizations of the training queries and sub-queries as well as the corresponding DB states and true cardinalities in the training part form the training data. We use the training data to train the query-driven models including ALECE, MLP, MSCN and NNGP. The data-driven models, namely DeepDB, NeuroCard, FLAT and FactorJoin, are built with the database data after all the statements in the training part of the workload are executed. This setting is compatible with real world scenarios as the cardinality estimation models are updated at regular intervals. In our experiments, the associated changing rate ρ of each testing query is larger than 20%. In other words, when testing queries are executed, at least 20% of the underlying data are changed, compared to the database when cardinality estimators are built.

The estimation results by different methods will be fed into our improved benchmark to compare the end-to-end query times (E2E times) of all methods on dynamic workloads. To investigate the performance gap between these methods and the optimal, we also feed the true cardinalities to the benchmark to get the optimal execution time. We do not compare with DeepDB on the STATS and TPCB datasets since it supports PK-FK joins only. The open implementation of FactorJoin is hard-coded for the STATS and Job-light datasets and does not support the TPCB dataset. Thus, FactorJoin is not tested on this dataset. Besides the E2E time, we also record the Q-error and P-error distributions, building time, storage

overhead and latency of all methods on different workloads. These results are shown in Table 4. The storage overhead, building time and latency results on the **Update-heavy** workloads are similar to the counterparts on the Insert-heavy workloads and thus omitted. **End-to-end evaluations.** Referring to Table 4, ALECE has clear advantages on the E2E time over the other methods. On the one hand, compared to PostgreSQL’s built-in cardinality estimator, ALECE makes query execution up to 2.7× faster. Its E2E time is only slightly larger than that of PG on TPCCH’s Insert-heavy workload. However, there is almost no gap between the performance of PG and Optimal on this workload. Compared to Optimal, ALECE only results in at most 2.2% extra E2E time. Considering the superiority of ALECE on these workloads, we confidently claim that ALECE performs much better than the built-in estimator in PostgreSQL and is able to greatly improve the query execution performance on dynamic workloads. On the other hand, in addition to ALECE, only MLP and Uni-Samp outperform PG on few workloads. The E2E time of DeepDB, NeuroCard, FLAT, FactorJoin and NNGP is larger than that of PG in all cases. Compared to Optimal, these methods results in more than double of E2E time on most workloads. These results indicate that these existing methods cannot make satisfactory estimates for DBMS on the dynamic workloads.

The P-error comparisons demonstrate that ALECE results in effective query executions from another angle. The P-errors of ALECE on most queries are close to 1. In other words, with ALECE’s outputs, most queries can be executed almost as fast as if they are optimized with true cardinalities. At the 95% quantile, ALECE outperforms PG, Uni-Samp, DeepDB, NeuroCard, FLAT, FactorJoin, MLP, MSCN and NNGP, in terms of P-error, by up to 26, 32, 21, 518, 1,385, 551, 1,443, 463 and 37 times, respectively. All these show that ALECE achieves large superiority over the competitors on helping PostgreSQL process queries more efficiently.

It is noteworthy that ALECE has clearer advantages on the STATS dataset with more complex join patterns. This reflects that our ALECE is able to grasp the implicit relations between complex join patterns and the underlying data significantly better than the alternatives. Besides, ALECE outperforms MLP in terms of E2E time and P-error. The main structure difference of ALECE and MLP is reflected on the adoption of attentions in ALECE’s modules. This verifies the positive effects of the attention mechanisms in extracting useful information from underlying data and SQL queries.

Estimation Accuracy. Table 4 also reports on the Q-error distributions of all methods. In general, ALECE clearly outperforms PG on all three datasets. This verifies that the independence assumption in PG is not reasonable for some scenarios. ALECE still performs best among all methods in all cases. At most quantiles, ALECE results in the smallest maximum Q-error. The median Q-error of ALECE in all cases are all close to 1, the optimal value. At the 95% quantile, the Q-error of ALECE in all cases is smaller than 10. In contrast, none of the other methods can reach this level of performance. Also at the 95% quantile, Uni-Samp, DeepDB, NeuroCard, FLAT, FactorJoin, MLP, MSCN and NNGP result in up to or more than $10^9\times$, $84\times$, $1,643\times$, $8,220\times$, $10^4\times$, $3,480\times$, $10^4\times$, $7,927\times$, $10^4\times$ larger Q-error than that of ALECE. At other quantiles, these methods are still incomparable with ALECE. These results show that ALECE is more accurate and able to discover the implicit relationships among attributes and those between queries and attributes. Another inter-

esting thing is that smaller Q-error does not necessarily result in smaller E2E time, which is claimed in the existing work [40] and shown by the comparisons among NNGP, DeepDB and NeuroCard, etc, on the Job-light dataset. An inaccurate estimate for a single sub-query tends to generate bad query plans and large execution time. It is necessary to ensure accurate estimates for all sub-queries. **Model Construction Efficiency.** Referring to Table 4, the training cost of ALECE is small. It requires less than 10, 8 and 13 minutes to fine-tune its parameters for the STATS, Job-light and TPCCH datasets, respectively. In contrast, DeepDB, NeuroCard, FLAT and MSCN consume more construction time. Although Uni-Samp and Factorjoin require less construction time in some cases, their E2E time, P-error and Q-error performance are much worse. Compared to ALECE, MLP and NNGP have simpler structures and thus they need less time to train. However, their representation abilities are not so powerful as that of ALECE. The overwhelming advantage of ALECE on E2E time illustrates the necessity of a more complex structure and slightly more training time.

In terms of latency, Uni-Samp achieves the worst performance. FactorJoin, MLP and MSCN require the least average time, less than 10 ms, on making an estimate. ALECE’s latency is slightly larger but smaller than the others. Considering the fact that executing a query on all datasets takes more than 10 seconds on average, estimation latency of less than 10 ms is not crucial in the overall picture.

Storage Overhead. The storage overhead of the query-driven methods is smaller than that of the data-driven ones in general. The query-driven methods only need to maintain a fixed number of parameters whose sizes are usually much smaller than the ‘data summaries’ held by the data-driven methods, e.g., the SPN or FSPN in DeepDB and FLAT, respectively. FactorJoin and MSCN incur the smallest memory costs. ALECE consumes more memory than MLP and NNGP due to its more complex structure. Compared to NeuroCard and FLAT, ALECE saves up to 75.4% and 89.4% memory cost on the STATS dataset, respectively. ALECE results in a little more memory cost than FLAT and DeepDB on the Job-light dataset. Considering the fact that modern computers usually have large memories and ALECE performs better in terms of E2E time, P-error and Q-error, the extra storage overhead pays off highly.

6.3 Effect of Distribution Shifting

To investigate if ALECE still works well when the distribution of the underlying data greatly changes, we carry out experiments to compare the different methods on the Dist-shift workload. It is noteworthy that the evaluation part of the Dist-shift workload covers highly skewed insert statements. Table 5 reports the E2E time, Q-error and P-error distribution comparisons among ALECE, PG, NeuroCard, MSCN, NNGP and Optimal. Those competitors are chosen because they have better overall performance on the Insert-heavy and Update-heavy workloads. Due to space limit, the storage overhead, building time and latency comparisons are omitted. These results are similar to the counterparts on the other two workloads.

Referring to Table 5, the overall Q-error and P-error performance of all methods on the Dist-shift workloads are worse than the counterparts on the Insert-heavy workloads. This is reasonable because compared to the data upon which these methods are built, the distribution of the underlying data is greatly shifted when testing queries are executed. Nevertheless, ALECE still achieves the best E2E time

Table 5: Performance of methods on Dist-shift workloads

Data	Model	E2E Time(S)	Q-error				P-error			
			50%	90%	95%	99%	50%	90%	95%	99%
STATS	PG	8,432	189	1.4·10 ⁵	1.1·10 ⁶	1.9·10 ⁷	2.60	25.50	42.65	300
	Uni-Samp	7,524	1.32	>10 ¹⁰	>10 ¹⁰	>10 ¹⁰	1.19	12.62	38.63	87.59
	NeuroCard	27,252	14.30	996	5,051	3.9·10 ⁵	2.08	25.69	108	3,318
	MSCN	26,697	20.08	2,802	15,576	4.9·10 ⁵	2.75	46.51	465	83,452
	NNGP	10,537	9.47	785	4,370	2.4·10 ⁵	1.45	7.19	25.19	3,318
	ALECE	6,876	1.40	5.46	11.75	118.36	1.07	1.38	1.71	11.18
	Optimal	6,770	1	1	1	1	1	1	1	1
Job-light	PG	5,608	1.58	6.12	14.22	76.30	1.10	1.32	1.48	2.89
	Uni-Samp	5,906	1.44	3.46	10.78	>10 ¹⁰	1.16	1.68	1.90	3.28
	NeuroCard	17,720	14.15	822	4,799	3.5·10 ⁵	1.81	12.40	37.25	144
	MSCN	25,157	28.54	1,592	5,261	99,178	2.06	15.01	27.77	84.35
	NNGP	11,698	4.6·10 ⁸	>10 ¹⁰	>10 ¹⁰	>10 ¹⁰	3.50	7.90	11.24	23.38
	ALECE	4,763	1.16	2.23	4.34	16.22	1.07	1.24	1.32	1.77
	Optimal	4,708	1	1	1	1	1	1	1	1
TPCH	PG	3,377	1.23	3.94	5.85	10.28	1.10	1.39	1.64	3.06
	Uni-Samp	>30,000	1.29	3.24	>10 ¹⁰	>10 ¹⁰	1.79	2.28	62.54	301
	NeuroCard	11,362	44.65	9,151	34,541	2.8·10 ⁵	1.60	15.10	25.94	50.68
	MSCN	7,062	41.62	6,571	28,837	2.4·10 ⁵	3.64	9.17	22.00	202
	NNGP	18,244	1.4·10 ⁹	>10 ¹⁰	>10 ¹⁰	>10 ¹⁰	5.54	46.30	59.12	99.12
	ALECE	3,230	1.18	2.66	4.46	11.23	1.07	1.23	1.29	1.57
	Optimal	3,226	1	1	1	1	1	1	1	1

on all datasets. It needs up to 18.5% less time than PG to execute all testing queries. Its E2E time on all three workloads is close to the optimal results and much smaller than that of other competitors. In terms of Q-error and P-error, ALECE also performs best. At the 95% quantile, the largest Q-error and P-error of ALECE are 11.75 and 1.71, respectively. In contrast, the other competitors' Q-error and P-error are at least 2.48 and 1.44 times larger, respectively. All these demonstrate that ALECE is less sensitive to the distribution shifting and able to make accurate estimates even when the distributions of the underlying data changes significantly.

7 RELATED WORK

Data-driven Cardinality Estimators. Data-driven methods aim to describe the underlying data with statistical or machine learning models. The simple yet efficient 1-D Histogram [47] is used in many well-known DBMS like PostgreSQL. It assumes all attributes are mutually independent and maintains a 1-D (cumulative) histogram for each attribute. To address the problem of unreasonable independence assumption, M-D Histogram based methods [12, 21, 43, 54] build multi-dimensional histograms to model attribute dependency. Although such methods improve the accuracy, the decomposition of the joint attributes is still lossy. Also, they hardly work for queries with complex joins. Sampling-based methods [9, 10, 32, 35, 45] address join queries but they risk high variance and sampling failure when the data distribution or query is complex. Bayesian network (BN) based methods [11, 17, 52] use a directed acyclic graph to model the dependence among attributes, assuming that each attribute is conditionally independent of the remaining attributes given its parents' distributions. BayesCard [60] revitalizes BN using probabilistic programming to improve its inference and model construction speed. Recently, machine learning techniques are adopted in data-driven methods. Deep autoregressive models are adopted in Naru [64] and NeuroCard [63] to decompose the joint distribution of attributes to a product of conditional distributions. DeepDB [26] is built upon Sum-Product Network (SPN) [42] which approximates the joint distribution using several local and simple PDFs. FLAT [67] improves SPN by adopting a factorize-split-sum-product network (FSPN) [62] to adaptively decompose the joint distribution according to the attribute dependence level.

Query-driven Cardinality Estimators. Such estimators focus on modeling the relationships between queries and their true cardinalities. The feedbacks of past queries are utilized to correct and self-tune histograms [8, 15, 28, 48] and update statistical summaries [49, 57]. LW-XGB and LW-NN [14] formulate the cardinality estimation as a regression problem and apply gradient boosted trees and neural networks for regression, respectively. UAE-Q [61] applies the deep auto-regression models and differentiable progressive sampling via the Gumbel-Softmax trick to learn hidden information from queries. The KDE-based join estimators [29] combine kernel density estimation (KDE) with a query-driven tuning mechanism to estimate multivariate probability distributions of a relation and cardinalities of joins. Fauce [38] and NNGP [65] assume a query's cardinality follows a Gaussian distribution and adopt Deep Ensemble [33] and neural network Gaussian process [34] to predict the distribution's mean and variance.

A few works also consider both data and SQL queries. Wu *et al.* [58] propose a unified deep autoregressive model utilizing both data as unsupervised information and query workload as supervised information. Kipf *et al.* [32] concatenate basic relation information and query features together and use a multi-set convolutional neural network to process them and make estimates. Negi *et al.* [41] propose techniques to build sample tables on the join keys and use neural networks to extract information from queries. However, these methods require either samples over joined attributes which results in high sampling overhead, or unrenovable featurization/samples over static data which are not applicable for dynamic database. Also, compared to ALECE, these methods do not explain the links among data, queries and true cardinalities. Besides, Negi's work [41] supports PK-FK joins only. Dutt *et al.* [14] use neural networks and tree-based ensembles to extract information from data and queries to solve the problem of selectivity estimation on a single relation. Thus, their approach does not support joins. In addition, Han *et al.* [22] propose an end-to-end evaluation benchmark for cardinality estimators. Our used benchmark is an improved version of it. Sun *et al.* [51] make a comprehensive comparison of the existing cardinality estimators.

8 CONCLUSION AND FUTURE WORK

In this work, we design ALECE, a versatile learned cardinality estimation model, that makes accurate and high-quality estimates for SQL queries. Based on two delicate methods to featurize the underlying database data and the SQL queries, respectively, ALECE adopts the attention mechanisms in its two modules to understand the implicit relations between data and queries. The self-attention layer in the data-encoder module figures out the links among all database attributes. The query-analyzer takes the input of the query featurization and the output of the data-encoder, and puts attention on the more important parts of the data. Extensive experimental results show that ALECE clearly outperforms the state-of-the-art alternatives in terms of multiple evaluation metrics.

For future work, it is interesting to extend ALECE to more general aggregate analytic queries by replacing COUNT(*) with other aggregate functions. Also, it is relevant to explore if better data and query featurization methods exist. Moreover, it makes sense to use other types of attention functions in ALECE.

REFERENCES

- [1] <https://github.com/pfl-cs/ALECE>.
- [2] <https://relational.fit.cvut.cz/dataset/Stats>.
- [3] <http://homepages.cwi.nl/~boncz/job/imdb.tgz>.
- [4] https://www.tpc.org/tpc_documents_current_versions/current_specifications5.asp.
- [5] Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer Normalization. *arXiv preprint abs/1607.06450* (2016). <http://arxiv.org/abs/1607.06450>
- [6] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *ICLR*.
- [7] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *NeurIPS*.
- [8] Nicolas Bruno, Surajit Chaudhuri, and Luis Gravano. 2001. STHoles: A Multidimensional Workload-Aware Histogram. In *SIGMOD*. 211–222.
- [9] Surajit Chaudhuri, Rajeev Motwani, and Vivek R. Narasayya. 1999. On Random Sampling over Joins. In *SIGMOD*. 263–274.
- [10] Yu Chen and Ke Yi. 2017. Two-Level Sampling for Join Size Estimation. In *SIGMOD*. 759–774.
- [11] C. K. Chow and C. N. Liu. 1968. Approximating discrete probability distributions with dependence trees. *IEEE Trans. Inf. Theory* 14, 3 (1968), 462–467.
- [12] Amol Deshpande, Minos N. Garofalakis, and Rajeev Rastogi. 2001. Independence is Good: Dependency-Based Histogram Synopses for High-Dimensional Data. In *SIGMOD*. 199–210.
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*. 4171–4186.
- [14] Anshuman Dutt, Chi Wang, Azade Nazi, Srikanth Kandula, Vivek R. Narasayya, and Surajit Chaudhuri. 2019. Selectivity Estimation for Range Predicates using Lightweight Models. *Proc. VLDB Endow.* 12, 9 (2019), 1044–1057.
- [15] Dennis Fuchs, Zhen He, and Byung Suk Lee. 2007. Compressed histograms with arbitrary bucket layouts for selectivity estimation. *Inf. Sci.* 177, 3 (2007), 680–702.
- [16] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional Sequence to Sequence Learning. In *ICML*. 1243–1252.
- [17] Lise Getoor, Benjamin Taskar, and Daphne Koller. 2001. Selectivity Estimation using Probabilistic Models. In *SIGMOD*. 461–472.
- [18] Goetz Graefe and William J. McKenna. 1993. The Volcano Optimizer Generator: Extensibility and Efficient Search. In *ICDE*. IEEE Computer Society, 209–218.
- [19] Karol Gregor, Ivo Danihelka, Andriy Mnih, Charles Blundell, and Daan Wierstra. 2014. Deep AutoRegressive Networks. In *ICML*, Vol. 32. 1242–1250.
- [20] PostgreSQL Global Development Group. 1996. PostgreSQL. <https://www.postgresql.org>. (1996). Accessed: 2022-10-28.
- [21] Dimitrios Gunopulos, George Kollios, Vassilis J. Tsotras, and Carlotta Domeniconi. 2000. Approximating Multi-Dimensional Aggregate Range Queries over Real Attributes. In *SIGMOD*. 463–474.
- [22] Yuxing Han, Ziniu Wu, Peizhi Wu, Rong Zhu, Jingyi Yang, Liang Wei Tan, Kai Zeng, Gao Cong, Yanzhao Qin, Andreas Pfadler, Zhengping Qian, Jingren Zhou, Jiangneng Li, and Bin Cui. 2021. Cardinality Estimation in DBMS: A Comprehensive Benchmark Evaluation. *Proc. VLDB Endow.* 15, 4 (2021), 752–765.
- [23] Trevor Hastie, Robert Tibshirani, and Jerome H. Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, 2nd Edition*. Springer.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. 770–778.
- [25] Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching Machines to Read and Comprehend. In *NeurIPS*. 1693–1701.
- [26] Benjamin Hilprecht, Andreas Schmidt, Moritz Kulessa, Alejandro Molina, Kristian Kersting, and Carsten Binnig. 2020. DeepDB: Learn from Data, not from Queries! *Proc. VLDB Endow.* 13, 7 (2020), 992–1005.
- [27] Raghav Kaushik, Jeffrey F. Naughton, Raghu Ramakrishnan, and Venkatesan T. Chakaravarthy. 2005. Synopses for query optimization: A space-complexity perspective. *ACM Trans. Database Syst.* 30, 4 (2005), 1102–1127.
- [28] Andranik Khachatryan, Emmanuel Muller, Christian Stier, and Klemens Bohm. 2015. Improving Accuracy and Robustness of Self-Tuning Histograms by Sub-space Clustering. *IEEE Trans. Knowl. Data Eng.* 27, 9 (2015), 2377–2389.
- [29] Martin Kiefer, Max Heimes, Sebastian Bref, and Volker Markl. 2017. Estimating Join Selectivities using Bandwidth-Optimized Kernel Density Models. *Proc. VLDB Endow.* 10, 13 (2017), 2085–2096.
- [30] Yoon Kim, Carl Denton, Luong Hoang, and Alexander M. Rush. 2017. Structured Attention Networks. In *ICLR*.
- [31] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [32] Andreas Kipf, Thomas Kipf, Bernhard Radke, Viktor Leis, Peter A. Boncz, and Alfons Kemper. 2019. Learned Cardinalities: Estimating Correlated Joins with Deep Learning. In *CIDR*.
- [33] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In *NIPS*. 6402–6413.
- [34] Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S. Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. 2018. Deep Neural Networks as Gaussian Processes. In *ICLR*.
- [35] Feifei Li, Bin Wu, Ke Yi, and Zhuoyue Zhao. 2016. Wander Join: Online Aggregation via Random Walks. In *SIGMOD*. 615–629.
- [36] Pengfei Li, Wenqing Wei, Rong Zhu, Bolin Ding, Jingren Zhou, and Hua Lu. 2023. ALECE: An Attention-based Learned Cardinality Estimator for SPJ Queries on Dynamic Workloads (Extended). *arXiv preprint abs/2310.05349* (2023). <https://arxiv.org/abs/2310.05349>
- [37] Xi Liang, Stavros Sintos, Zechao Shang, and Sanjay Krishnan. 2021. Combining Aggregation and Sampling (Nearly) Optimally for Approximate Query Processing. In *SIGMOD*. 1129–1141.
- [38] Jie Liu, Wenqian Dong, Dong Li, and Qingqing Zhou. 2021. Fauce: Fast and Accurate Deep Ensembles with Uncertainty for Cardinality Estimation. *Proc. VLDB Endow.* 14, 11 (2021), 1950–1963.
- [39] Guido Moerkotte, Thomas Neumann, and Gabriele Steidl. 2009. Preventing Bad Plans by Bounding the Impact of Cardinality Estimation Errors. *Proc. VLDB Endow.* 2, 1 (2009), 982–993.
- [40] Parimarjan Negi, Ryan C. Marcus, Andreas Kipf, Hongzi Mao, Nesime Tatbul, Tim Kraska, and Mohammad Alizadeh. 2021. Flow-Loss: Learning Cardinality Estimates That Matter. *Proc. VLDB Endow.* 14, 11 (2021), 2019–2032.
- [41] Parimarjan Negi, Ziniu Wu, Andreas Kipf, Nesime Tatbul, Ryan Marcus, Sam Madden, Tim Kraska, and Mohammad Alizadeh. 2023. Robust Query Driven Cardinality Estimation under Changing Workloads. *Proc. VLDB Endow.* 16, 6 (2023), 1520–1533.
- [42] Hoifung Poon and Pedro M. Domingos. 2011. Sum-Product Networks: A New Deep Architecture. In *UAI*. 337–346.
- [43] Viswanath Poolsala and Yannis E. Ioannidis. 1997. Selectivity Estimation Without the Attribute Value Independence Assumption. In *VLDB*. 486–495.
- [44] Lutz Prechelt. 2012. Early Stopping - But When? In *Neural Networks: Tricks of the Trade - Second Edition*. Lecture Notes in Computer Science, Vol. 7700. Springer, 53–67.
- [45] Yuan Qiu, Yilei Wang, Ke Yi, Feifei Li, Bin Wu, and Chaoqun Zhan. 2021. Weighted Distinct Sampling: Cardinality Estimation for SPJ Queries. In *SIGMOD*. 1465–1477.
- [46] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. 1986. *Learning Internal Representations by Error Propagation*. 318–362.
- [47] Patricia G. Selinger, Morton M. Astrahan, Donald D. Chamberlin, Raymond A. Lorie, and Thomas G. Price. 1979. Access Path Selection in a Relational Database Management System. In *SIGMOD*. 23–34.
- [48] Utkarsh Srivastava, Peter J. Haas, Volker Markl, Marcel Kutsch, and Tam Minh Tran. 2006. ISOMER: Consistent Histogram Construction Using Query Feedback. In *ICDE*. 39.
- [49] Michael Stillger, Guy M. Lohman, Volker Markl, and Mokhtar Kandil. 2001. LEO - DB2’s LEarning Optimizer. In *VLDB*. 19–28.
- [50] Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. End-To-End Memory Networks. In *NeurIPS*. 2440–2448.
- [51] Ji Sun, Jintao Zhang, Zhaoyan Sun, Guoliang Li, and Nan Tang. 2021. Learned Cardinality Estimation: A Design Space Exploration and A Comparative Evaluation. *Proc. VLDB Endow.* 15, 1 (2021), 85–97.
- [52] Kostas Tzoumas, Amol Deshpande, and Christian S. Jensen. 2011. Lightweight Graphical Models for Selectivity Estimation Without Independence Assumptions. *Proc. VLDB Endow.* 4, 11 (2011), 852–863.
- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NeurIPS*. 5998–6008.
- [54] Hai Wang and Kenneth C. Sevcik. 2003. A multi-dimensional histogram for selectivity estimation and fast approximate query answering. In *CASCON*. 328–342.
- [55] Jiayi Wang, Chengliang Chai, Jiabin Liu, and Guoliang Li. 2021. FACE: A Normalizing Flow based Cardinality Estimator. *Proc. VLDB Endow.* 15, 1 (2021), 72–84.
- [56] Xiaoying Wang, Changbo Qu, Weiyuan Wu, Jiannan Wang, and Qingqing Zhou. 2021. Are We Ready For Learned Cardinality Estimation? *Proc. VLDB Endow.* 14, 9 (2021), 1640–1654.
- [57] Chenggang Wu, Alekh Jindal, Saeed Amizadeh, Hiren Patel, Wangchao Le, Shi Qiao, and Sriram Rao. 2018. Towards a Learning Optimizer for Shared Clouds. *Proc. VLDB Endow.* 12, 3 (2018), 210–222.
- [58] Peizhi Wu and Gao Cong. 2021. A Unified Deep Model of Learning from both Data and Queries for Cardinality Estimation. In *SIGMOD*. 2009–2022.

- [59] Ziniu Wu, Parimarjan Negi, Mohammad Alizadeh, Tim Kraska, and Samuel Madden. 2023. FactorJoin: A New Cardinality Estimation Framework for Join Queries. *Proc. ACM Manag. Data* 1, 1 (2023), 41:1–41:27.
- [60] Ziniu Wu and Amir Shaikhha. 2020. BayesCard: A Unified Bayesian Framework for Cardinality Estimation. *CoRR* abs/2012.14743 (2020). <https://arxiv.org/abs/2012.14743>
- [61] Ziniu Wu, Pei Yu, Peilun Yang, Rong Zhu, Yuxing Han, Yaliang Li, Defu Lian, Kai Zeng, and Jingren Zhou. 2022. A Unified Transferable Model for ML-Enhanced DBMS. In *CIDR*.
- [62] Ziniu Wu, Rong Zhu, Andreas Pfadler, Yuxing Han, Jiangneng Li, Zhengping Qian, Kai Zeng, and Jingren Zhou. 2020. FSPN: A New Class of Probabilistic Graphical Model. *CoRR* abs/2011.09020 (2020). <https://arxiv.org/abs/2011.09020>
- [63] Zongheng Yang, Amog Kamsetty, Sifei Luan, Eric Liang, Yan Duan, Xi Chen, and Ion Stoica. 2020. NeuroCard: One Cardinality Estimator for All Tables. *Proc. VLDB Endow.* 14, 1 (2020), 61–73.
- [64] Zongheng Yang, Eric Liang, Amog Kamsetty, Chenggang Wu, Yan Duan, Xi Chen, Pieter Abbeel, Joseph M. Hellerstein, Sanjay Krishnan, and Ion Stoica. 2019. Deep Unsupervised Cardinality Estimation. *Proc. VLDB Endow.* 13, 3 (2019), 279–292.
- [65] Kangfei Zhao, Jeffrey Xu Yu, Zongyan He, Rui Li, and Hao Zhang. 2022. Lightweight and Accurate Cardinality Estimation by Neural Network Gaussian Process. In *SIGMOD*. 973–987.
- [66] Zhuoyue Zhao, Robert Christensen, Feifei Li, Xiao Hu, and Ke Yi. 2018. Random Sampling over Joins Revisited. In *SIGMOD*. 1525–1539.
- [67] Rong Zhu, Ziniu Wu, Yuxing Han, Kai Zeng, Andreas Pfadler, Zhengping Qian, Jingren Zhou, and Bin Cui. 2021. FLAT: Fast, Lightweight and Accurate Method for Cardinality Estimation. *Proc. VLDB Endow.* 14, 9 (2021), 1489–1502.