

NeuroCard: One Cardinality Estimator for All Tables

Zongheng Yang
UC Berkeley
zongheng@berkeley.edu

Amog Kamsetty*
UC Berkeley
amogkamsetty@berkeley.edu

Sifei Luan*
UC Berkeley
lsf@berkeley.edu

Eric Liang
UC Berkeley
ericliang@berkeley.edu

Yan Duan
Xi Chen
Covariant
{rocky,peter}@covariant.ai

Ion Stoica
UC Berkeley
istoica@berkeley.edu

ABSTRACT

Query optimizers rely on accurate cardinality estimates to produce good execution plans. Despite decades of research, existing cardinality estimators are inaccurate for complex queries, due to making lossy modeling assumptions and not capturing inter-table correlations. In this work, we show that it is possible to learn the correlations across all tables in a database without any independence assumptions. We present NeuroCard, a join cardinality estimator that builds a single neural density estimator over an entire database. Leveraging join sampling and modern deep autoregressive models, NeuroCard makes no inter-table or inter-column independence assumptions in its probabilistic modeling. NeuroCard achieves orders of magnitude higher accuracy than the best prior methods (a new state-of-the-art result of $8.5\times$ maximum error on JOB-light), scales to dozens of tables, while being compact in space (several MBs) and efficient to construct or update (seconds to minutes).

PVLDB Reference Format:

Zongheng Yang, Amog Kamsetty, Sifei Luan, Eric Liang, Yan Duan, Xi Chen, and Ion Stoica. NeuroCard: One Cardinality Estimator for All Tables. PVLDB, 14(1): 61 - 73, 2021.
doi:10.14778/3421424.3421432

PVLDB Artifact Availability:

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

1 INTRODUCTION

Query optimizers translate queries into executable plans with the best estimated performance. They are critical not only for relational databases, but also for modern analytics engines, such as Spark [1] and Presto [36]. Among various techniques, *cardinality estimation* often plays a larger role than the cost model or the plan search space in producing high-quality query plans [19]. Unfortunately, cardinality estimation is a notoriously difficult problem, where the accuracy may drop exponentially as the query complexity (e.g., the number of joins) increases [21].

*Equal contribution.

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. 14, No. 1 ISSN 2150-8097.
doi:10.14778/3421424.3421432

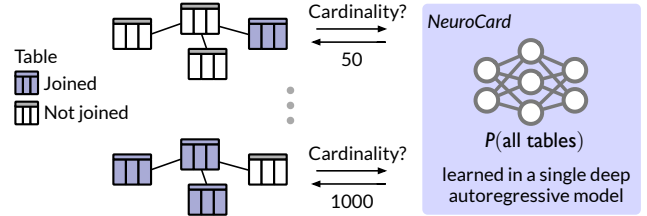


Figure 1: NeuroCard uses a single probabilistic model, which learns all possible correlations among all tables in a database, to estimate join queries on any subset of tables.

At a high level, there are two approaches to cardinality estimation: *query-driven* and *data-driven*. Query-driven estimators typically rely on supervised learning to learn a function mapping (featurized) queries to predicted cardinalities. They implicitly assume queries from a production workload are “similar” to training queries—namely, training and test sets of queries are drawn from the same underlying distribution. This assumption can be violated when, for example, users issue unexpected types of queries.

In contrast, data-driven estimators approximate the data distribution of a table—a function mapping each tuple to its probability of occurrence in the table—instead of training on “representative” queries. A simple method to approximate the data distribution is a histogram. In theory, once we estimate the distribution of each table in a schema, we can estimate the output cardinality of any query. While this approach is more general, it suffers from two drawbacks: (1) *lossy modeling assumptions* (e.g., assume the tables’ distributions are independent), and (2) *low precision* (e.g., a limited number of histogram bins). Fortunately, recent advances in machine learning have alleviated both drawbacks. Unlike previous density estimators, *deep autoregressive (AR) models* [4, 6, 32, 33, 42] can learn complex high-dimensional data distributions without independence assumptions, achieving state-of-the-art results in both precision and expressiveness. This has resulted in new data-driven cardinality estimators based on deep AR models [48].

However, despite their promise, deep autoregressive model-based cardinality estimators are limited to handling single tables. There are three challenges that make this approach ineffective for *joins*:

- **High training cost:** To learn the distribution of a join, any data-driven estimator needs to see actual tuples from the join result. Unfortunately, for all but the smallest scale, it is expensive, and sometimes infeasible, to precompute the join.
- **Lack of generality:** The AR approach builds a probabilistic model for each join, e.g., $T = T_1 \bowtie T_2 \bowtie T_3$, that it estimates.

However, the model for T cannot be directly used to estimate a join on a subset of T , e.g., $T_2 \bowtie \sigma(T_3)$. Of course, one could train a model for every possible join. This can be prohibitive, as the number of possible joins is exponential in the number of tables.

- **Large model size:** The complexity of the learned AR model grows with the cardinality of the dataset. As joins tend to involve columns with high cardinalities, an AR model built on a join may incur a prohibitively large size.

In this work, we propose NeuroCard, a learning-based join cardinality estimator that directly learns from data to overcome these challenges. NeuroCard’s distinctive feature is the ability to capture the correlations across multiple joins in a single deep AR model, without any independence assumptions (Figure 1). Once trained, this model can handle all queries issuable to the schema, regardless of what subset of tables is involved. We address the above challenges using the following key ingredients.

To reduce training cost, NeuroCard samples from a join, instead of computing the join fully (§4). The key property of such a sample is to capture the join’s distribution: if a key is more frequent in the join result, it should be more frequent in the sample as well. To meet this requirement, we precompute the correct sampling weights for each key. While the worst-case cost of computing the join is exponential in the number of tables, computing the sampling weights is done in time linear with the data size by dynamic programming.

To achieve generality, NeuroCard needs to train a single model to answer queries on any subset of tables (§6). The basic idea behind our solution is to train the AR model on samples from the *full outer join* of all tables. The full join contains the values from all the base tables, so it has sufficient information to answer a query touching any subset of tables. At inference time, if a table in the schema is not present in a join query, we need to account for any potential *fanout* effect. Consider an AR model trained on samples from the full join $T = T_1 \bowtie T_2$, and a query $\sigma(T_1)$ whose cardinality we want to estimate. If the join key of T_2 is the foreign key of T_1 , then a tuple of T_1 may *appear multiple times* in T . NeuroCard learns the probabilities of these “duplicated” tuples and additional bookkeeping information, which enables us to account for fanouts.

Finally, to scale to large-cardinality columns while avoiding prohibitively large models, NeuroCard employs *lossless column factorization* (§5). An AR model stores one embedding vector per distinct value, so it could quickly blow up in size for columns with large numbers of distinct values, e.g., 100,000s or more. With factorization, a column is decomposed into several subcolumns, each taking a chunk of bits from the binary representation of the original column values. For instance, a 32-bit ID column id can be decomposed into (id_0, \dots, id_3) with the first subcolumn corresponding to the first 8 bits, and so on. We then train the autoregressive model on these lower-cardinality subcolumns instead of the full columns.

By combining these ingredients, NeuroCard achieves state-of-the-art estimation accuracy, including in the challenging tail quantiles. On the popular JOB-light benchmark, a schema that contains 6 tables and basic filters, NeuroCard achieves a maximum Q-error of 8.5× using 4 MB. This corresponds to a 4.6× improvement over the previous state of the art. We created a more difficult benchmark, JOB-light-ranges, with a larger variety of content columns and range filters. On this benchmark, NeuroCard achieves up to 15–34×

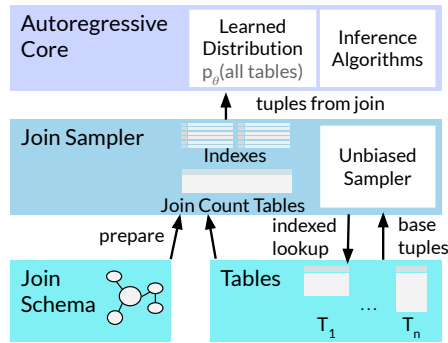


Figure 2: Overview of NeuroCard. The Join Sampler (§4) provides correct training data (sampled tuples from join) by using unbiased join counts. Sampled tuples are streamed to an autoregressive model for maximum likelihood training (§3). Inference algorithms (§6) use the learned distribution to estimate query cardinalities.

higher accuracy than previous solutions, including DeepDB [12], MSCN [15], and IBJS [20]. Lastly, to test NeuroCard’s ability to handle a more complex join schema, we created JOB-M which has 16 tables and multi-key joins. NeuroCard scales well to this benchmark, offering 10× higher accuracy than conventional approaches while maintaining a low model size (27 MB, covering 16 tables).

In summary, this paper makes the following contributions:

- We design and implement NeuroCard, the first learned data-driven cardinality estimator that learns across joins in a schema without any independence assumptions. All in-schema correlations among the tables are captured by a single autoregressive model, which can estimate any query on any subset of tables.
- NeuroCard learns the correct distribution of a join without actually computing that join. Instead, the model is trained on uniform and independent samples of the join of all tables in a schema.
- We propose lossless column factorization (§5), a technique that significantly reduces the size of the autoregressive model, making its use practical for high-cardinality columns.
- Compared to the best prior methods, NeuroCard significantly improves the state-of-the-art accuracy on the JOB-light benchmark. We further propose two new benchmarks, JOB-light-ranges and JOB-M, and show that both are much more challenging and thus better gauges of estimator quality (§7).

To invite further research, NeuroCard and the benchmarks used in this paper are open source at <https://github.com/neurocard>.

2 OVERVIEW OF NEUROCARD

Consider a set of tables, T_1, \dots, T_N . We define their *join schema* as the graph of join relationships, where vertices are tables, and each edge connects two joinable tables. A query is a subgraph of the overall schema. If a query joins a table multiple times, our framework duplicates that table in the schema. We assume the schema and queries submitted to the estimator are acyclic (§4.2 discusses relaxations), so they can be viewed as trees.

Next, we present an overview of NeuroCard as a sequence of goals and solutions to achieve these goals.

2.1 Goals and Solutions

Goal: A single estimator. Our goal is building a single cardinality estimator for the entire join schema. For example, assuming the schema has three tables, the estimator should handle joins on any subset of tables, e.g., $\sigma(T_2)$, $T_1 \bowtie T_3$, or $T_1 \bowtie T_2 \bowtie \sigma(T_3)$.

Having a single estimator has two key benefits: simplicity and accuracy. Having multiple estimators—each covering a specific join template (a table subset)—does not scale for a large number of tables, as the number of possible join templates increases exponentially. In addition, it is easier for a DBMS to operationalize a single estimator rather than many estimators. Most importantly, having multiple estimators can hurt accuracy. This is because estimating the cardinality of a query on a table subset not covered by any single estimator, but by multiple estimators, requires some form of independence assumption to combine these estimators. If the independence assumption does not hold, the accuracy will suffer.

Solution: We build a single cardinality estimator that learns the distribution of *the full outer join of all tables* in the schema (henceforth, full join). For example, for a three-table schema, we learn $p(T_1 \bowtie T_2 \bowtie T_3)$. Note that using the inner join instead of the full join would not work. Indeed, the inner join $T_1 \bowtie T_2 \bowtie T_3$ is the intersection of the three tables. If a query uses only T_1 or $T_1 \bowtie T_3$, their tuples may not be fully contained in this intersection, and thus the estimator would have insufficient information to answer these queries.

Goal: Efficient sampling of the full join. A data-driven estimator learns a distribution by reading representative tuples from that distribution. To learn the distribution of the full join, a straightforward approach is to compute it and then uniformly draw random samples from the result. Unfortunately, even on a small 6-table schema (the JOB-light workload), the full join contains *two trillion* ($2 \cdot 10^{12}$) tuples, making it infeasible to compute in practice.

Solution: We perform uniform sampling over the full join *without* materializing it. Specifically, we ensure that any tuple in the full join J (a multiset) is sampled with same probability, $1/|J|$. To achieve this, we leverage a state-of-the-art join sampling algorithm [50] (§4). We first precompute *join count tables* that map each table’s join keys to their correct sampling weights with respect to the full join. Then, we sample the keys using these counts as weights. Given a sampled key, we construct the full tuple by looking up the remaining columns via indexes¹ from all tables, and then concatenating them. This way, we only need to materialize the join counts as opposed to the full join. Using dynamic programming, computing the join counts takes time linear in the size of the database, and is quite fast in practice (e.g., 13 seconds for 6 tables in JOB-light, and 4 minutes for 16 tables in JOB-M).

Goal: Support any subset of tables. Although the full outer join contains all information of the tables, we need to take care when a query involves just a subset of the tables. Consider:

$$T_1.\text{id} : [1, 2] \quad T_2.\text{id} : [1, 1] \quad \longrightarrow \quad T_1 \bowtie T_2 : [(1, 1), (1, 1), (2, \emptyset)]$$

Query: $\sigma_{\text{id}=1}(T_1)$

¹Like prior work on join sampling [20, 22], we assume base tables have an index built for each join key. This impacts the efficiency but not correctness of the design.

The correct selectivity is $\frac{1}{2}$ (1 row). However, in the full join distribution, $P(T_1.\text{id} = 1) = \frac{2}{3}$ (2 rows). This is because we have not accounted for the *fanout* produced by the missing table, T_2 .

Solution: Handle *schema subsetting*: If a query does not include a table, we downscale the estimate by the fanout introduced by that table. In essence, since the learned probability space is the full join, we must downscale appropriately when a query touches a subset and expects the returned selectivity to refer to that subset.

Goal: Accurate density estimation. The final ingredient to achieve our goal is an accurate and compact density estimator.

Solution: We leverage *deep autoregressive (AR) models* to implement our density estimator. This family of neural density estimators have been successfully employed on high-dimensional data types such as image [33], audio [42], and text [32]. Recently, Naru [48] has leveraged deep AR models to achieve state-of-the-art accuracy results on estimating the cardinalities of single-table queries, while *learning the correlations among all columns* without independence assumptions. We apply Naru to learn the distribution of the full join, and optimize its construction and inference for our setting.

2.2 Putting It All Together

Figure 2 shows the high-level architecture of NeuroCard.

Building the estimator consists of two stages. First, we prepare the join sampler by building or loading existing single-table indexes on join keys and computing the join count tables for the specified join schema (§4). Second, we train the deep AR model by repeatedly requesting batches of sampled tuples from the sampler, usually 2K tuples at a time. The sampler fulfills this request in the background, potentially using multiple sampling threads.

Once the estimator is built, it is ready to compute the cardinality estimates for given queries. For each query, we use probabilistic inference algorithms (§6) to compute the cardinality estimate by (1) performing Monte Carlo integration on the learned AR model, and (2) handling schema subsetting. A single estimator can handle queries joining any subset of tables, with arbitrary range selections.

3 CONSTRUCTING NEUROCARD

In this section, we present the background of the techniques used to implement NeuroCard.

3.1 Probabilistic Modeling of Tables

Consider a table T with column domains $\{A_1, \dots, A_n\}$. This table induces a discrete *joint data distribution*, defined as the probability of occurrence of each tuple ($f(\cdot)$ denotes number of occurrences):

$$p(a_1, \dots, a_n) = f(a_1, \dots, a_n)/|T|.$$

The n -dimensional data distribution (the *joint*) $p(\cdot)$ allows us to compute a query’s cardinality as follows. Define a query Q as $\sigma : A_1 \times \dots \times A_n \rightarrow \{0, 1\}$. Then, the *selectivity*—the fraction of records that satisfy the query—can be computed as a probability: $P(Q) = \sum_{a_1 \in A_1} \dots \sum_{a_n \in A_n} \sigma(a_1, \dots, a_n) \cdot p(a_1, \dots, a_n)$. The *cardinality* is obtained by multiplying it with the row count: $|Q| = P(Q) \cdot |T|$.

Data-driven cardinality estimators can be grouped along two axes: (1) joint factorization, and (2) the density estimator used.

Joint factorization, or the modeling assumption, determines how precisely data distribution p is factored. Any modeling assumption risks losing information about correlations across columns, which ultimately leads to a loss in accuracy. For example, the widely used ID histogram technique assumes the columns are independent. As a result, it factors p into a set 1D marginals, $p \approx \prod_{i=1}^n p(A_i)$, which can lead to large inaccuracies when the columns’ values are strongly correlated. Similarly, other data-driven cardinality estimators such as graphical models [3, 7, 8, 40, 41] either assume conditional independence or partial independence among columns. One exception is the autoregressive (product-rule) factorization,

$$p = \prod_{i=1}^n p(A_i | A_{<i}), \quad (1)$$

which precisely expresses the overall joint distribution as the product of the n conditional distributions.

The density estimator determines how precisely the aforementioned factors are actually approximated. The most accurate “estimator” would be recording these factors exactly in a hash table. Unfortunately, this leads to enormous construction and inference costs (e.g., in the case of $p(A_n | A_{1:n-1})$). At the other end, the 1D histogram has low costs, but this comes at the expense of low precision, as it makes no distinction between the values falling in the same bin. Over the years, a plethora of solutions have been proposed, including kernel density estimators and Bayesian networks. Recently, *deep autoregressive (AR) models* [4, 32, 33] have emerged as the density estimator of choice. Deep AR models compute $\{p(A_i | A_{<i})\}$ without explicitly materializing them by learning the n conditional distributions in compact neural networks. Deep AR models achieve state-of-the-art precision, and, for the first time, provide a tractable solution for implementing the autoregressive factorization.

3.2 Naru: Deep Autoregressive Models as Cardinality Estimators

NeuroCard builds on Naru, a state-of-the-art cardinality estimator that fully captures the correlations among all columns of a single table using a deep AR model. Next, we present an overview of Naru and discuss how NeuroCard leverages it.

Construction. Given table T , an AR model θ takes a tuple $\mathbf{x} \in T$ as input, and predicts conditional probability distributions, $\{p_\theta(X_i | \mathbf{x}_{<i})\}$, each of which is an 1D distribution over the i -th column (conditioned on all prior column values of \mathbf{x}). The likelihood of the input tuple is then predicted as $p_\theta(\mathbf{x}) = \prod_{i=1}^n p_\theta(X_i = \mathbf{x}_i | \mathbf{x}_{<i})$. Any deep AR architecture can instantiate this framework, e.g., ResMADE [4] or Transformer [43]. Training aims to approximate the data distribution p using p_θ , by minimizing the KL divergence [26], $D_{KL}(p || p_\theta)$. This is achieved by maximum likelihood estimation (MLE) and gradient ascent to maximize the predicted (log-)likelihood of data:

$$\text{Sample i.i.d. } \mathbf{x} \sim p \quad (2)$$

$$\text{Take gradient steps to maximize } \log p_\theta(\mathbf{x}) \quad (3)$$

In our setting, we define T as the full outer join of all tables within a schema. Consequently, the deep AR model learns the correlations across all tables. Next, we need to sample tuples with probabilities prescribed by p . Otherwise, p_θ would approximate an incorrect,

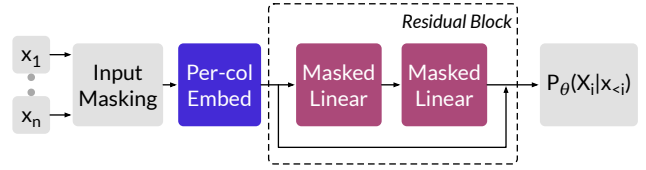


Figure 3: Default architecture of the autoregressive model.

biased distribution. To achieve this, we use a sampler that emits *simple random samples* from the full join T (§4).

Estimating query cardinalities. Once constructed, the Naru estimator estimates the cardinality of a given query. A query is represented as a hyper-rectangle: each column X_i with domain A_i is constrained to take on values in a valid region $R_i \subseteq A_i$:

$$\text{Query: } \bigwedge \{X_i \in R_i\} \quad (4)$$

Next, Naru estimates the probability of the query (an event) using a Monte Carlo integration algorithm, *progressive sampling*:

$$\text{ProgressiveSampling}(\{X_i \in R_i\}): p_\theta(\bigwedge \{X_i \in R_i\}) \cdot |T| \quad (5)$$

It works by drawing imaginary, in-region tuples from the model’s learned distributions. Specifically, it draws the first dimension of the sample as $x_1 \sim p_\theta(X_1 | X_1 \in R_1)$, the second dimension of the sample as $x_2 \sim p_\theta(X_2 | X_2 \in R_2; x_1)$, and so on. The likelihoods of the samples are importance-weighted. This procedure also efficiently supports omitted columns, i.e., wildcards of the form $X_i \in *$.

NeuroCard’s inference invokes progressive sampling to estimate cardinalities, but extends it in two ways. First, we apply the column factorization optimization (§5), which potentially changes a X_{i+1} ’s valid region, R_{i+1} , based on the value drawn from X_i . Second, we add support for schema subsetting (§6), by downscaling selectivity $p_\theta(\bigwedge \{X_i \in R_i\})$ by the corresponding fanout.

3.3 Join Problem Formulation

A join schema induces the full outer join of all tables in the schema, $T = T_1 \bowtie \dots \bowtie T_N$. Our goal is to build a fully autoregressive probabilistic model on the full join consisting of all tables’ columns:

$$\text{Model: } p_\theta(T) \equiv p_\theta(T_1.\text{col}_1, T_1.\text{col}_2, \dots, T_N.\text{col}_k) \quad (6)$$

We can then use the probabilistic model to estimate the cardinalities of join queries on any subset of tables in the schema.

Supported joins. NeuroCard supports acyclic join schemas and queries containing multi-way, multi-key equi-joins (§4.2 discusses how to relax the acyclic requirement). The schema should capture the most common joins. For joins not captured in the schema, their cardinalities can be estimated by first obtaining single-table estimates using NeuroCard, then combining the estimates using classical heuristics [21]. This allows uncommon cases to be handled under the same framework, albeit at the cost of lower accuracy.

Supported filters. NeuroCard supports equality and range filters on discrete or numerical columns. These include arithmetic comparison operators ($<$, $>$, \leq , \geq , $=$) and IN. More complex filters can also be expressed using the valid region encoding, mentioned in the previous section. Arbitrary forms of AND/OR can be handled via the inclusion-exclusion principle.

3.4 Model architecture

NeuroCard uses a standard AR architecture, ResMADE [4], which is also employed by Naru; see Figure 3. Input tuples are represented as discrete, dictionary-encoded IDs, (x_1, \dots, x_n) , and embedded by per-column embedding matrices. The concatenated embedded vector is fed to a series of residual blocks, each consisting of two masked linear layers (they are masked to ensure the autoregressive property). The output layer produces logits $\{\log p_\theta(X_i | x_{<i})\}$ by dotting the last layer’s output with the embedding matrices. Next, we compute a cross-entropy loss on the logits and perform back-propagation. We turn on Naru’s *wildcard skipping* optimization, which randomly masks inputs to train special marginalization tokens that aid infer-time estimation (i.e., using these tokens to skip sampling any wildcards in a query).

Masked multi-layer perceptrons such as ResMADE strike a good balance between efficiency and accuracy. NeuroCard can use any advanced AR architectures, if desired. In §7, we also instantiate NeuroCard with an advanced architecture (the Transformer [43]).

4 SAMPLING FROM JOINS

A key challenge in NeuroCard is computing an *unbiased* sample of the full join (§2.1) to ensure that the learned distribution faithfully approximates the full join distribution. Namely, every tuple in the full join J (a multiset) must be sampled equally likely with probability $1/|J|$. The samples should also be i.i.d., as required by Equation 2. NeuroCard meets these requirements by using a sampler that produces *simple random samples with replacement*.

4.1 Algorithm

A tuple in the full join contains *join key columns* and *content columns*. Our sampler exploits this decomposition. The first step of the sampler is to precompute *join count tables*, which are per-table statistics that reflect the occurrence counts of the join keys in the full join. The sampler then samples the join keys, table-by-table, with occurrence probabilities proportional to their join counts. Lastly, it selects content columns from the base tables by looking up the drawn join keys. This completes a batch of sample, which is sent to the model for training, and the procedure repeats on demand.

Computing join counts. Zhao *et al.* [50] provide an efficient algorithmic framework of join sampling that produces simple random samples from general multi-key joins. NeuroCard implements the Exact Weight algorithm from Zhao *et al.*, adapted to full outer joins.

We illustrate the algorithm on a join schema (a tree) consisting of tables T_1, \dots, T_N . For exposition, assume they only involve join keys (content columns are gathered later). Let T_1 be the root table. The join count of a tuple $t \in T_i$ is the total number of tuples in the full outer join of all of T_i ’s descendants that joins with t . It is recursively defined as:

$$w_i(t) = \prod_{T_j \in \text{Children}(T_i)} \sum_{t' \in t \times T_j} w_j(t') \quad \forall i, \forall t \in T_i \quad (7)$$

where $t \times T_j$ denotes all tuples in T_j that join with t . For a leaf table with no descendants, $w_i(\cdot)$ is defined as 1. At the root table T_1 , $w_1(t)$ represents the count of all $t \in T_1$ in the entire full outer join. The join counts of each table are computed by aggregating over the join counts of all of its child tables, and can thus be computed

recursively in a bottom-up fashion. Using dynamic programming, the time complexity is linear in the number of tuples in all tables, $O(|T_1| + \dots + |T_N|)$.

Sampling. Once the join counts are computed, the sampler produces a sample by traversing the join tree in a top-down fashion. It starts by drawing a sample t_1 from the root table T_1 using weights $\{w_1(t) : t \in T_1\}$ (i.e., with probabilities $\{w_1(t)/\sum_{t' \in T_1} w_1(t')\}$). It then samples through all descendants of T_1 in the breadth-first order. At a child table, say T_2 , it samples t_2 from $t_1 \times T_2$ (all tuples in T_2 that join with t_1) using weights $\{w_2(t) : t \in t_1 \times T_2\}$. The procedure continues recursively until all tables are visited, and thus produces a sample (t_1, \dots, t_N) , each t_i being a tuple of join keys from the respective table.

Example. Consider the schema in Figure 4a. Figure 4b shows the computed join counts. The leaf table C has a count of 1 for every tuple. In B , since $(2, c)$ can join with two tuples in C , its join count is $2 = 1 + 1$. Similar propagation happens for $A.x = 2$ which gets a count of $3 = 1 + 2$. Physically, we store the join counts indexed by join keys (e.g., for C , only one mapping $c \rightarrow 1$ is kept). For sampling, suppose $A.x = 2$ is first sampled. It has two matches in B with weights 1 and 2, so the second match, $(2, c)$, has an inclusion probability of $2/3$.

NULL handling. To support full outer joins, we handle NULL keys as follows. We add a virtual \emptyset tuple (which denotes NULL) to each table T_i , and make it join with all normal $t \in T_j$ that have no matches in T_i , where $T_j \in \text{Children}(T_i)$. Similarly, any normal $t \in \text{Parent}(T_i)$ that has no match in T_i joins with T_i ’s \emptyset . All-NULL is invalid. Propagation proceeds as before; Figure 4b shows examples.

Constructing complete sample tuples. In the prior example, suppose $\langle 2; 2, c; c \rangle$ is drawn. We gather the content columns of A by looking up $A.x = 2$ and similarly for $(B.x, B.y) = (2, c)^2$ and $C.y = c$. On multiple matches, we pick a row uniformly at random. Their concatenation represents a sampled tuple from the full join.

Computing the size of the full join (normalizing constant). Recall from §3.2 that the row count $|J|$ (the *normalizing constant* in probabilistic terms) is required to convert selectivities into cardinalities. With join counts it can be computed exactly: $|J| = \sum_{t \in T_1} w_1(t)$.

Parallel sampling. Finally, the sampling procedure is embarrassingly parallel: after the join count tables $\{w_i(\cdot)\}$ are produced, parallel threads can be launched to read the join counts and produce samples. Computation of the join count tables is also parallelizable, although it is an one-time effort. Sampling correctness is preserved even in the presence of parallelism due to the i.i.d. property.

4.2 Comparison with other samplers

Our key requirements of uniform and i.i.d. samples from the full join render many related sampling algorithms unsuitable. If either property is not satisfied, the sampling distribution would be biased and thus compromise the quality of the learned AR model. As examples, Index-based Join Sampling (IBJS) [20] is neither uniform nor independent; Wander Join [22] produces independent but non-uniform samples. Both approaches do produce unbiased

²Either intersect two matching lists from both columns’ index lookups, or do a single lookup if a composite index is available.

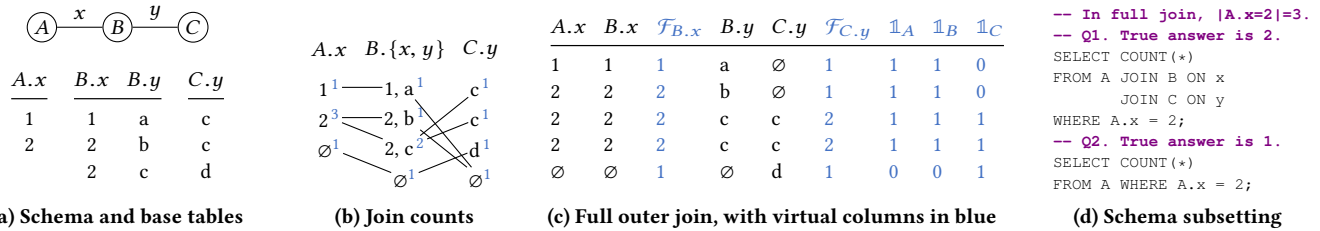


Figure 4: End-to-end example. (a) A join schema of three tables and their join key columns. Content columns are omitted. (b) Join counts (blue) enable uniform sampling of the full outer join and are computed in linear time by dynamic programming. Here, edges connect join partners. (c) Learning target: the full outer join of the schema, with *virtual columns* in blue. We show the *fanouts* \mathcal{F} , the number of times a join key value appears in the corresponding base table, for keys $B.x$ and $C.y$. The fanouts for $A.x$ and $B.y$ are all 1 and omitted. Each *indicator* $\mathbb{1}_T$ denotes whether a tuple has a match in table T . (d) Examples of schema subsetting, i.e., queries that touch a subset of the full join (§6).

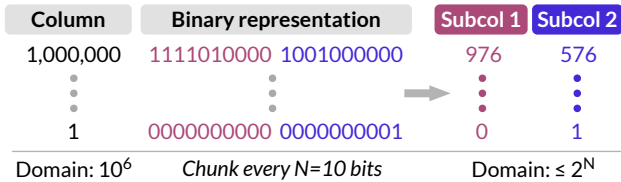


Figure 5: Lossless column factorization (§5).

estimators for counts or other aggregate statistics, but are not designed to return uniform join samples. Reservoir sampling, a well-known technique, draws samples without replacement (thus, non-independent) and requires a full scan over the full join, which is not scalable. Lastly, the Exact Weight algorithm NeuroCard implements is among the most efficient in Zhao *et al.* [50]. They provide additional extensions to support general, potentially cyclic joins (e.g., a cycle can be *broken*), which NeuroCard can leverage to broaden our formulation (§3.3).

5 LOSSLESS COLUMN FACTORIZATION

A key challenge of using an autoregressive model for high-cardinality data is that the size of the model parameters can scale linearly with the numbers of distinct values in the columns. In the model architecture we use (§3.4), each column (any data type; categorical or numerical) is first dictionary-encoded into integer token IDs. Then a per-column *embedding* layer is applied on these token IDs. The size of the trainable embedding matrix (essentially, a hash table) for each column C scales linearly with $|C|$, i.e., the number of distinct values in the domain. Even a moderately sized column with up to 10^6 distinct values, therefore, easily takes up 128 MB of space, assuming 32-dimensional embeddings are used.

To handle high-cardinality columns efficiently, we propose an optimization that we call *lossless column factorization*. This optimization is inspired by the popular use of “subword units” [35] in modern natural language processing, and also shares characteristics with “bit slicing” in the indexing literature [28]. Different from subword units, column factorization does not use a statistical algorithm such as byte pair encoding to determine what subwords to use (a potential optimization). Different from bit slicing, we slice a value into groups of bits and convert them back into base-10 integers.

Figure 5 illustrates the idea on a simple example. Suppose a column (any datatype) has a domain size of $|C| = 10^6$. Naively supporting this column would require allocating $|C| \cdot h$ floats as

its embedding matrix, where h is the embedding dimension. Instead, NeuroCard factorizes each value *on-the-fly* during training: we convert an original-space value into its binary representation, then slice off every N bits, the *factorization bits* hyperparameter. Each sliced off portion becomes a *subcolumn*, now in base-10 integer representation. These subcolumns are now treated as regular columns to learn over by the autoregressive model. Crucially, a much smaller embedding matrix is now needed for each subcolumn containing at most $2^N \cdot h$ floats. In this example, we can reduce 128 MB to 250 KB—a more than 500× space reduction.

Model size vs. statistical efficiency. Choosing the factorization bits N enables a tradeoff between model size vs. statistical efficiency. By decreasing N , we have more subcolumns, each with a smaller domain, but learning across more variables becomes harder. In theory, by using autoregressive modeling no information is lost in this translation, so the precision of the learned distributions is not affected. In practice, we observed that lower factorization bits, i.e., slicing into more subcolumns, generally underperform higher ones that use more space, but not by a significant margin (§7.5). We thus set the factorization bits N based on a space usage budget.

Lossless = factorization + autoregressive modeling. With factorization, a column is factorized into multiple subcolumns, which are then fed into a downstream density estimator. However, if a density estimator with independence assumptions, e.g., 1D histograms, is used, then this whole process is *lossy*. By modeling $p(\text{subcol}_1, \text{subcol}_2) \approx p(\text{subcol}_1)p(\text{subcol}_2)$, histograms would fail to capture any potential correlation between the two subcolumns. In other words, other estimators *could* read in subcolumn values and potentially reduce space usage, but their inherent quality and assumptions determine how much information is learned about the subcolumns, and about their correlations with other columns. By using autoregressive modeling, NeuroCard forces the AR model to explicitly capture such correlation, namely (ignoring other columns):

$$p(\text{col}) \equiv p(\text{subcol}_1, \text{subcol}_2) = p(\text{subcol}_1)p(\text{subcol}_2|\text{subcol}_1),$$

which has no inherent loss of information. Hence, we call the unique combination of factorization and autoregressive modeling *lossless*.

Filters on subcolumns. During probabilistic inference, a filter on an original column needs to be translated into *equivalent* filters on subcolumns. Recall from §3.2 that the probabilistic inference

procedure draws samples that lie inside the queried region. We modify that procedure to handle subcolumns by respecting each filter’s semantics. Going back to our example, consider the filter $\text{col} < 1,000,000$. The filter for the high-bits subcol_1 is *relaxed* to ≤ 976 (note the less-equal). The inference procedure would draw a subcol_1 value in this range, based on which the low-bits filter is relaxed appropriately. If the drawn subcol_1 is 976, then the filter on subcol_2 is set to “ < 576 ”; otherwise, the high-bits already satisfy the original filter so a wildcard is placed on the low-bits subcolumn. This is reminiscent of processing range predicates on bit-sliced indexes [28]; NeuroCard applies these processing logic in the new context of probabilistic inference for autoregressive models.

6 QUERYING NEUROCARD

Once built, the autoregressive model summarizes the entire full outer join. The challenge with querying this probabilistic model for a selectivity estimate is that the query may *restrict the space it touches to a subset of the full join*—a phenomenon we term *schema subsetting*. Since the selectivity estimate returned by the model assumes the probability space to be the full outer join, rather than the query-specific restricted space, the estimate should be *downscaled* appropriately during probabilistic inference.

NeuroCard’s inference algorithms combine two building blocks. First, Naru [48] introduced *progressive sampling*, a Monte Carlo algorithm that integrates over an autoregressive model to produce selectivity estimates. We invoke this routine (i.e., Equation 5) on the trained autoregressive model with changes outlined in this section. Second, Hilprecht *et al.* [12] have proposed inference algorithms to query a sum-product network trained on a full outer join. We state their algorithms below and discuss how to adapt these algorithms into our framework, thereby generalizing them to a new type of probabilistic model.

Basic case: no table omitted. The simplest case of schema subsetting is an *inner* join query on all tables. Consider the example data in Figure 4a and an inner join query Q1 in Figure 4d. The query, $\sigma_{A.x=2}(A \bowtie_x B \bowtie_y C)$, restricts the probability space from the full join to the inner join. Naively querying the model for $|A.x = 2|$ would return a cardinality of $|J| \cdot (3/5) = 3$ rows, as 3 out of 5 rows in the full join J (Figure 4c) satisfy the filter. However, the correct row count for this query is 2 (two rows in the inner join; both pass the filter). Left/right outer joins can also exhibit this behavior.

To correct for this, Hilprecht *et al.* propose a simple solution by adding an *indicator column* per table into the full join. A binary column $\mathbb{1}_T$ is added for each table T , with value 1 if a tuple (in the full join) has a non-trivial join partner with table T , and 0 otherwise.

NeuroCard adopts this solution as follows. First, during training, the sampler is tasked with appending these *virtual* indicator columns on-the-fly to sampled tuples. Recall that each sampled tuple is formed by querying base-table indexes with sampled join keys. If a table T contains a join key, we set that sampled tuple’s $\mathbb{1}_T$ to 1, and 0 otherwise (see Figure 4c). The autoregressive model treats these indicator columns as regular columns to be learned.

Second, during inference, NeuroCard adds equality constraints on the indicator columns, based on what tables are present in the query. The progressive sampling routine (Equation 5) not only gets the usual filter conditions, $\{X_i \in R_i\}$, but also $\{\mathbb{1}_T = 1\}$ for any

table T that appears in the inner-join query graph³. In summary, for the no-omission case, the routine now estimates the probability:

$$P(\{X_i \in R_i\} \wedge \{\mathbb{1}_T = 1 : \text{for all table } T\}) \quad (8)$$

Example. Coming back to the example query Q1, $\sigma_{A.x=2}(A \bowtie_x B \bowtie_y C)$, we compute the selectivity under the full join as $P(A.x = 2 \wedge \mathbb{1}_A = \mathbb{1}_B = \mathbb{1}_C = 1)$. Reading from Figure 4c, this probability is $2/5$, so the cardinality is correctly computed as $5 \cdot (2/5) = 2$ rows.

Omitting tables and fanout scaling. The less straightforward case is if a query *omits*, i.e., does not join, certain tables. Consider Q2 in Figure 4d: $\sigma_{A.x=2}(A)$. When restricting the scope to table A , the row count of $A.x = 2$ is 1, different from $|J| \cdot P(A.x = 2 \wedge \mathbb{1}_A = 1) = 3$ rows. The fundamental reason this happens is because the operation of a full join has *fanned out* tuples from base tables. To correctly downscale, Hilprecht *et al.* propose recording a per-join *fanout* column. We adapt this solution in NeuroCard⁴.

Specifically, for each join key column $T.k$, we insert into the full join a virtual fanout column, $\mathcal{F}_{T.k}$, defined as the number of times each value appears in $T.k$. For example, 2 appears twice in $B.x$, so its fanout is $\mathcal{F}_{B.x}(2) = 2$; see Figures 4a and 4c. Again, we task the join sampler with adding these fanout values on-the-fly to each batch of sampled tuples. The inclusion of fanouts is piggybacked onto the index lookup path (querying the size of each lookup result list), which adds negligible overheads.

On the inference side, Hilprecht *et al.* showed that the correct cardinality with omitted tables can be computed via *fanout scaling*:

$$\begin{aligned} \text{Cardinality}(\text{query } Q) &= |J| \cdot P(\{X_i \in R_i\} \text{ subsetted to query } Q) \\ &= |J| \cdot \mathbb{E}_{X \sim J} \left[\frac{\mathbb{1}_{\{X_i \in R_i\}} \cdot \prod_{T \in Q} \mathbb{1}_T}{\prod_{R \notin Q} \mathcal{F}_{R.\text{key}}} \right]. \end{aligned} \quad (9)$$

In essence, the numerator handles the basic case above, while the denominator counts the total number of times omitted tables $\{R \notin Q\}$ have fanned out each tuple in query Q . It loops through each omitted table R , finds its unique join key $R.\text{key}$ that connects to Q in the schema (discussed in detail below), and looks up the associated fanout value $\mathcal{F}_{R.\text{key}}$. We incorporate this scaling as follows. Since the fanout columns are learned by the model, we modify progressive sampling to draw a concrete value for each relevant $\mathcal{F}_{R.\text{key}}$ per progressive sample, compute the product of these fanouts, and divide the progressive sample’s estimated likelihood by this product.

Example. Coming back to Q2, $\sigma_{A.x=2}(A)$, the constraints are $\{A.x = 2, \mathbb{1}_A = 1\}$. Reading from Figure 4c, three rows satisfy the constraints and the relevant downscaling keys are $B.x$ and $C.y$. Thus the expectation expands as: $\frac{1}{5} \cdot (\frac{1}{2-1} + \frac{1}{2-2} + \frac{1}{2-2}) = \frac{1}{5}$. Multiplying with $|J| = 5$ arrives at the correct cardinality of 1 row.

Handling fanout scaling for multi-key joins. Our formulation of fanout scaling supports multi-key joins, e.g., both x and y keys in the example schema $A.x = B.x \wedge B.y = C.y$ (Figure 4a). The challenge of fanout scaling in this case is determining the set of omitted keys to downscale. Let V be the set of all tables. Let Q be

³The indicator columns can also be constrained appropriately for left or right joins.

⁴Our definition differs slightly from Hilprecht *et al.* In that work, each fanout column is bound to a PK-FK join and stores the frequency of a value in the FK. Our treatment binds a fanout to each join key, regardless of PK/FK, and is defined as the frequency each value appears in that key column itself. This removes their assumption of PK-FK joins and supports general equi-joins where both join keys can have duplicate values.

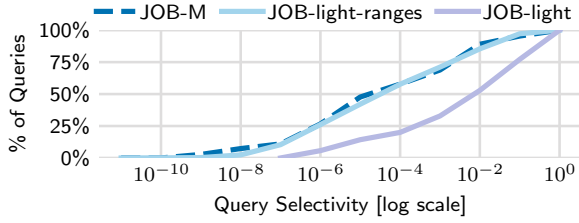


Figure 6: Distribution of query selectivity (§7.1).

the set of tables joined in a query, and the complement $O = V \setminus Q$ the omitted tables. Pick any table $T \in Q$. There exists a unique path from each omitted $T_O \in O$ to T , because the join schema graph is a tree (acyclic, connected). The join key attached to the edge incident to T_O on this path is the unique join key for table T_O to downscale. Hence, the fanout downscaling factor in Equation 9 is well-defined.

Going back to example Q2 where only A is queried, when considering the omitted table B which has two join keys $(B.x, B.y)$, we see that $B.x$ is the unique fanout key since it lies on the path $A \longleftrightarrow B$.

Summary of schema subsetting. To recap, NeuroCard’s probabilistic inference leverages the progressive sampling algorithm from Naru and the idea of additional columns from Hilprecht *et al.* that we term *virtual columns*. Our join sampler is modified to logically insert into the full join two types of virtual columns, the indicators and the fanouts. Both are treated as regular columns to be learned over by the density model, and both are used during progressive sampling to handle various cases of schema subsetting.

Ordering virtual columns in the autoregressive factorization. The autoregressive model requires some fixed ordering of columns in its factorization (§3.2). Naru has shown that different orderings may have different performance in the tail error but not in the lower error quantiles. We adopt the same practice as Naru in using an arbitrary ordering for the content columns. For the virtual columns introduced above, we place them after all the content columns, with indicators before fanouts. The intuition here is to ensure that (1) the conditional distributions involving content columns do not get confused by the presence of virtual columns, and (2) when sampling fanouts, placing them at the end allows for prediction using a maximum amount of prior information.

In our early benchmarks this choice performed better than if virtual columns were placed early in the ordering. We also experimented with *multi-order training* [6] in the autoregressive model, but did not see noticeably better performance. Thus, we opt for a simple treatment and leave such optimizations to future work.

7 EVALUATION

We evaluate NeuroCard on accuracy and efficiency and compare it with state-of-the-art cardinality estimators. The key takeaways are:

- **NeuroCard outperforms the best prior methods by 4–34× in accuracy (§7.3).** On the popular JOB-light benchmark, NeuroCard achieves a maximum error of 8.5× using 4 MB.
- **NeuroCard scales well to more complex queries (§7.3).** On the two new benchmarks JOB-light-ranges (more difficult range filters) and JOB-M (more tables in schema), NeuroCard achieves orders of magnitude higher accuracy than prior approaches.

Table 1: Workloads used in evaluation. **Tables:** number of base tables. **Rows, Cols, Dom.:** row count, column count, and maximum column domain size of the full outer join of each schema. **Feature characterizes each workload’s queries.** Rows in full join: $2 \cdot 10^{12}$; $2 \cdot 10^{12}$; 10^{13} .

WORKLOAD	TABLES	ROWS	COLS	DOM.	FEATURE
JOB-light	6	$2 \cdot 10^{12}$	8	235K	single-key joins
JOB-light-ranges	6	$2 \cdot 10^{12}$	13	134K	+complex filters
JOB-M	16	10^{13}	16	2.7M	+multi-key joins

- **NeuroCard is efficient to construct and query (§7.4).** A few million tuples, learned in less than 5 minutes, suffice for it to reach best-in-class accuracy.
- **We study the relative importance of each component of NeuroCard (§7.5).** Out of all factors, learning the correlations across all tables and performing unbiased join sampling prove the most impactful.

7.1 Experimental Setup

Workloads (Table 1). We adopt the real-world IMDB dataset and schema to test cardinality estimation accuracy. Prior work [19, 21] reported that correlations abound in this dataset and established it to be a good testbed for cardinality estimators. We test the following query workloads on IMDB:

- **JOB-light:** a 70-query benchmark used by many recent cardinality estimator proposals [12, 15, 38]. The schema contains 6 tables, title (primary), cast_info, movie_companies, movie_info, movie_keyword, movie_info_idx and is a typical star schema—every non-primary table only joins with title on title.id. The full outer join contains $2 \cdot 10^{12}$ tuples. Each query joins between 2 to 5 tables, with only equality filters except for range filters on title.production_year.
- **JOB-light-ranges:** we synthesized this second benchmark containing 1000 queries derived from JOB-light by enriching filter variety. We generate the 1000 queries uniformly distributed to each join graph of JOB-light (18 in total), as follows. For each join graph, using our sampler we draw a tuple from the inner join result. We use the non-null column values of this tuple as filter literals, and randomly place 3–6 comparison operators associated with these literals, based on whether each column can support range (draw one of $\{\leq, \geq, =\}$) or equality filters ($=$). Overall, this generator (1) follows the data distribution and guarantees non-empty results, and (2) includes more filters, in variety and in quantity, than JOB-light. An example 3-table query is: $mc \bowtie \sigma_{\text{info_type_id}=99}(\text{mi_idx}) \bowtie \sigma_{\text{episode_nr} \leq 4 \wedge \text{phonetic_code} \geq \text{N612}}(t)$, where t.id is joined with other tables’ movie_id.
- **JOB-M:** this last benchmark contains 16 tables in IMDB and involves *multiple* join keys. For instance, the table movie_companies is joined not only with title on movie_id, but also with company_name on company_id, and with company_type on company_type_id, etc. We adapt the 113 JOB queries [19] by allowing each table to appear at most once per query and removing logical disjunctions (e.g., $A.x=1 \vee B.y=1$). Each query joins 2–11 tables. We use JOB-M to test NeuroCard’s scalability as its full join is 5× larger and has more dimensions than the above (see Table 1). The benchmarks are available at <https://github.com/neurocard>.

Metric. We report the usual Q-error distribution of each workload, where the Q-error of a query is the multiplicative factor an estimated cardinality deviates from the query’s true cardinality: $Q\text{-error}(\text{query}) := \max\left(\frac{\text{card}_{\text{actual}}}{\text{card}_{\text{estimate}}}, \frac{\text{card}_{\text{estimate}}}{\text{card}_{\text{actual}}}\right)$. Both actual and estimated cardinalities are lower bounded by 1, so the minimum attainable Q-error is $1\times$. As reported in prior work [48], reducing high-quantile errors is much more challenging than mean or median; thus, we report the quantiles p_{100}, p_{99}, p_{95} , and the median. For timing experiments, we report latency/throughput using an AWS EC2 VM with a NVIDIA V100 GPU and 32 vCPUs.

Benchmark characteristics. Figure 6 plots the distributions of selectivities of these workloads, where we calculate each query’s selectivity as $\text{card}_{\text{actual}}/\text{card}_{\text{inner}}$ (denominator is the row count of the query join graph—an inner join—without filters). The selectivity spectrums of our two benchmarks (JOB-light-ranges and JOB-M) are much wider than JOB-light due to higher filter variety. The median selectivity is more than $100\times$ lower, while at the low tail the minimum selectivities are $1000\times$ lower.

7.2 Compared Approaches

We compare against several prevalent families of estimators. In each family, we aim to choose a state-of-the-art representative. Related Work (§8) includes a more complete discussion on all families and their representative methods.

Supervised query-driven estimators. We use MSCN [15] as a recent representative from this family. It takes in a featurized query, runs the query filters on pre-materialized samples of the base tables, then use these bitmaps as additional network inputs, and predicts a final cardinality. For JOB-light, we used the training queries and sample bitmaps provided in the authors’ source code [16]. For JOB-light-ranges, due to new columns, we generated 10K new training queries—generating and executing them to obtain true cardinality labels took 3.2 hours—and used a bitmap size of 2K to match the size of other estimators in this benchmark. For JOB-light, we also cite the best numbers obtained by Sun and Li [38], termed *E2E*, which is a deep supervised net with more effective building blocks (e.g., pooling, LSTM) than MSCN.

Unsupervised data-driven estimators. We use DeepDB [12] as a recent technique in this family. It uses a non-neural sum-product network [29] as the density estimator for each table subset chosen by correlation tests. Conditional independence is assumed across subsets. In contrast, NeuroCard uses a neural autoregressive model to build a single learned estimator over all tables in a schema. We use two recommended configurations from DeepDB: a base version that learns up to 2-table joins, and a larger version that additionally builds 3-table models. These correspond to their storage-optimized and the standard setups, respectively.

We found that the DeepDB source code [13] did not support range queries on categorical string columns out-of-the-box. Since JOB-light-ranges contains such queries, we perform data and query rewriting for this baseline, by dictionary-encoding the string values into integers. Reported results are with this optimization enabled.

Join sampling. We implement the Index-based Join Sampling method (IBJS) [20], using 10,000 as the maximum sample size. A

Table 2: JOB-light, estimation errors. Lowest errors are bolded.

ESTIMATOR	SIZE	Median	95th	99th	Max
Postgres	70 KB	7.97	797	$3 \cdot 10^3$	10^3
IBJS	–	1.48	10^3	10^3	10^4
MSCN	2.7 MB	3.01	136	$1 \cdot 10^3$	10^3
E2E (quoting [38])	N/A	3.51	139	244	272
DeepDB	3.7 MB	1.32	4.90	33.7	72.0
DeepDB-large	32 MB	1.19	4.66	35.0	39.5
NeuroCard	3.8 MB	1.57	5.91	8.48	8.51

query’s cardinality is estimated by taking a sample from the query’s join graph and executing per-table filters on-the-fly.

Real DBMS. We use Postgres v12, which performs cardinality estimation using 1D histograms and heuristics to combine them.

Other baselines. The methods chosen above have been compared to other estimators in prior studies. Naru [48] has shown that estimators based on classical density modeling (KDE; Bayesian networks; the MaxDiff n-dimensional histogram) or random sampling significantly lag behind deep autoregressive models. DeepDB [12] also shows that it significantly outperforms wavelets [2]. We therefore do not compare to these methods.

NeuroCard. We implement NeuroCard on top of the Naru source code [27]. We use ResMADE by default. For complex benchmarks we also use the Transformer (§3.4), which is suffixed with -large.

7.3 Estimation Accuracy

7.3.1 JOB-light. Table 2 reports each estimator’s accuracy on the 70 JOB-light queries. Overall, NeuroCard exhibits high accuracy across the spectrum. **It sets a new state-of-the-art maximum error at $8.5\times$** using 3.8 MB of parameters. This represents an $> 8\times$ improvement over the best prior method when controlling for size.

We now discuss a few observations. Not surprisingly, Postgres has the most inaccurate median—indicating a systematic mismatch between the approximated distribution and data—due to its use of coarse-grained density models (histograms) and heuristics. IBJS fares better at the median, but falls off sharply at tail, because samples of a practical size have a small chance to hit low-density queries in a large joint space. Both MSCN and E2E are deep supervised regressors which show marked improvements over prior methods. However, their median and 95th errors are quite similar and have sizable gaps from the two data-driven estimators.

NeuroCard vs. DeepDB shows interesting trends. NeuroCard is up to $4\text{--}8\times$ better at tail (99th, max), and DeepDB is slightly better at lower quantiles. NeuroCard is more robust at tail due to (1) a markedly better density model (neural autoregressive vs. non-neural sum-product networks that use inter-column independence assumptions), and (2) learning all possible correlations among the columns of all 6 tables, whereas DeepDB assumes (conditional) independence across several table subsets. DeepDB-large, being $8.4\times$ bigger and trained on $7.7\times$ more (54M) tuples, still trails NeuroCard at tail by more than $4\times$. NeuroCard slightly trails at the lower quantiles (“easy” queries with high true density) likely due to the mode-covering behavior of KL-divergence minimization [9].

Table 3: JOB-light-ranges, estimation errors. Lowest errors bolded.

ESTIMATOR	SIZE	Median	95th	99th	Max
Postgres	70 KB	13.8	$2 \cdot 10^3$	$2 \cdot 10^4$	$5 \cdot 10^6$
IBJS	–	10.1	$4 \cdot 10^4$	10^6	10^8
MSCN	4.5 MB	4.53	397	$6 \cdot 10^3$	$2 \cdot 10^4$
DeepDB	4.4 MB	3.40	537	$8 \cdot 10^3$	$2 \cdot 10^5$
DeepDB-large	21 MB	2.00	91.7	$2 \cdot 10^3$	$4 \cdot 10^4$
NeuroCard	4.1 MB	1.87	57.1	375	8169
NeuroCard-large	21 MB	1.40	35.1	232	1029

Table 4: JOB-M, estimation errors. Lowest errors are bolded.

ESTIMATOR	SIZE	Median	95th	99th	Max
Postgres	120 KB	174	$1 \cdot 10^4$	$8 \cdot 10^4$	$1 \cdot 10^5$
IBJS	–	61.1	$3 \cdot 10^5$	$4 \cdot 10^6$	$4 \cdot 10^6$
NeuroCard	27.3 MB	2.84	404	1327	$2 \cdot 10^4$
NeuroCard-large	409 MB	1.96	26.4	304	874

7.3.2 *JOB-light-ranges*. This 1000-query benchmark adds equality/range filters on more content columns, using the same join templates as JOB-light (which has range filters on one column only). Results are shown in Table 3.

NeuroCard achieves the best accuracy across all error quantiles, and improves on the best prior methods by up to 15–34×. It is also the only estimator with $< 2\times$ median and 3-digit 99%-tile errors. Overall, all estimators produce less accurate cardinalities, though the drops are of varying degrees. Compared with MSCN, NeuroCard improves by 2× at median, 7× at 95th, 15× at 99th, and 2× at max. Compared with DeepDB, NeuroCard improves the four quantiles by 2×, 9×, 21×, and 23×, respectively. Comparing the enlarged versions of the two estimators (suffixed with -large), the accuracy gains become 1.4×, 2.6×, 9.6× and 34×, respectively.

NeuroCard’s improvements over baselines significantly widen in this benchmark, due to prior approaches failing to capture the more complex inter-column correlations being tested.

7.3.3 *JOB-M*. This final benchmark tests NeuroCard’s ability to scale to a much larger and more complex join schema. Different from the JOB-light schema, JOB-M contains 16 tables, with each query joining 2–11 tables on multiple join keys (in addition to movie_id only in JOB-light). For baselines, we only include Postgres and IBJS, because MSCN’s query encoding does not support the complex filters in this benchmark and DeepDB ran out of memory on this 16-table dataset due to high-cardinality categorical columns.

Results in Table 4 show that **NeuroCard’s accuracy remains high on this complex schema**. Postgres produces large errors, and IBJS also struggles, due to many intermediate samples becoming empty as the number of joins grows. NeuroCard overcomes this challenge and offers more than 10× better accuracy across the board. In terms of space efficiency, since the model needs to be trained on the full outer join of 16 tables and the maximum domain size exceeds 2 million, a vanilla NeuroCard would require 900 MB in model size. With column factorization (§5), the model size is reduced to 27MB—less than 1% of the total size of all tables. We also present a large model NeuroCard-large to demonstrate scalability.



Figure 7: Statistical and physical efficiency of NeuroCard.

7.4 Efficiency

Having established that NeuroCard achieves the best accuracy, we now study the statistical and physical efficiency of NeuroCard.

How many tuples are required for good accuracy? Figure 7a plots accuracy (p99 on JOB-light and JOB-light-ranges) vs. number of tuples trained. *About 2–3M tuples are sufficient for NeuroCard to achieve best-in-class accuracy* (compare with Tables 2 and 3). Using more samples helps, but eventually yields diminishing returns. Reaching high accuracy using a total of $\sim 10^7$ samples out of a population of 10^{12} data points (i.e., only 0.001% of the data)—many queries would inevitably touch unseen data points—shows that NeuroCard generalizes well and is *statistically efficient*.

How does sampling affect training throughput? Figure 7b plots the training throughput, in tuples per second, vs. the number of sampling threads used to provide training data. Four threads suffice to saturate the GPU used for training. At lower thread counts, the device spends more time waiting for training data than doing computation. With a peak throughput of $\sim 40K$ tuples/second, NeuroCard can finish training on 3M tuples in about 1.25 minutes.

Wall-clock training time comparison. Figure 7c compares the wall-clock time used for training the MSCN, DeepDB, and NeuroCard configurations reported in Tables 2 and 3. MSCN requires a separate phase of executing training queries to collect true cardinalities, which takes much longer (3.2 hours for 10K queries) than just the training time shown here. DeepDB runs on parallel CPUs and is quite efficient. NeuroCard starts training/on-the-fly sampling after calculating the join count tables, which takes 13 seconds for both datasets. Its construction is efficient due to parallel sampling and accelerated GPU computation.

Wall-clock inference time comparison. Lastly, Figure 7d plots the latency CDF of the learning approaches for 1000 JOB-light-ranges queries. As before, we use the base configurations reported in the accuracy Tables. MSCN and NeuroCard run on GPU while DeepDB runs on CPU; all three approaches are implemented in Python. MSCN is fastest because its lightweight network has fewer

Table 5: Ablation studies: varying primary components of NeuroCard. Unlisted values are identical to the Base configuration. We show the impact of the sampler (A), column factorization bits (B), autoregressive model size (C), inter-table correlations learned (D), and whether to use an autoregressive model at all (E) on the 50% and 95%-tile errors of JOB-light-ranges.

	Sampler	Fact. Bits	$d_{ff}; d_{emb}$	Correlations Learned	p50	p95
Base (4.1 MB)	unbiased	14	128; 16	all tables in one AR	1.9	57.1
(A)	biased				33	3270
(B)		10 (2.2 MB)			2.2	173
		12 (2.6 MB)			2.0	168
		None (12 MB)			1.6	62.7
(C)			128; 64 (23 MB)		1.5	44.0
			1024; 16 (31 MB)		1.7	64.0
(D)				one AR per table	40	$9 \cdot 10^4$
(E)	No model; uniform join samples only				4.0	$2 \cdot 10^5$

calculations involved. DeepDB’s latencies span a wide spectrum, from ~ 1 ms for queries with low complexity (numbers of joins and filters involved) to ~ 100 ms for queries with the highest complexity. NeuroCard’s latencies are more predictable, with 17 ms at median and 12 ms at minimum: this is due to the higher number of floating point operations involved in the neural autoregressive model. All approaches can be sped up by engineering efforts (e.g., if run in a native language). For NeuroCard, model compression or weight quantization can also reduce the computational cost.

7.5 Dissecting NeuroCard

To gain insights, we now evaluate the relative importance of primary components of NeuroCard, by varying them and measuring the change in estimation accuracy on JOB-light-ranges. We use the smaller NeuroCard in Table 3 as the *Base* configuration, and ablate each component in isolation. Table 5 presents the results.

In (A), using IBJS adapted for full joins⁵ as a *biased* sampler significantly decreases the learned estimator’s accuracy. The large increase in the median error implies a systematic distribution mismatch. Overall, this design choice is the second most important.

Rows in group (B) vary the column factorization granularity. Using smaller bits results in more subcolumns and yields a small drop in accuracy. Disabling factorization uses the most space and appears to perform the best.

Group (C) varies the size of the autoregressive model, by changing the dimension of the feedforward linear layers (d_{ff}) or the embeddings (d_{emb}). An enlarged embedding proves markedly more useful than enlarged linear layers, likely because each token’s captured semantics becomes more finetuned during optimization.

⁵The fact table title is ordered at front and a large intermediate size of 10^6 is used.

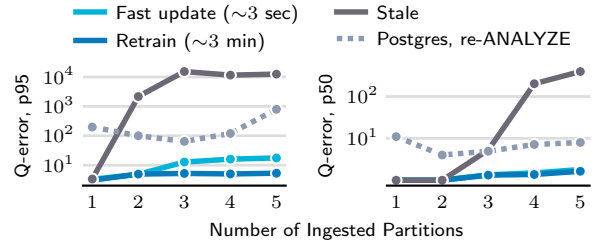


Figure 8: Updating NeuroCard, fast and slow. JOB-light. Errors (p95, p50) of each strategy are averaged from 10 runs. Postgres is also run as comparison, whose statistics are updated (1~2 sec.) on each ingest.

In group (D) we vary the correlation learned by NeuroCard. While all configurations above learn the distribution of all tables in a single model—capturing all possible correlations among them—here we build one model (same architecture as Base) per table. Queries that join across tables are estimated by combining individual models’ estimates via independence. Without modeling inter-table correlations, this variant yields the lowest accuracy.

Finally, group (E) ablates away the AR model altogether. We test *uniform join samples* as a standalone estimator: it uses our sampler (§4) to draw 10^4 simple random samples (actual tuples in the database) from each query’s join graph. While the median error is reasonable, it is $10^4\times$ less accurate than an autoregressive model at tail as many queries have no sample hits. The AR model is more statistically efficient than sampling, because it provides access to conditional probability distributions—these conditional contributions enable an efficient probabilistic inference procedure, i.e., progressive sampling, which cannot be used otherwise.

Tuning guide. Groups (B) and (C) show that NeuroCard is not overly sensitive to hyperparameters. For new datasets, we recommend starting with the Base configuration and increasing sizing as much as possible up to a size budget. The recommended precedence is: factorization bits; d_{emb} ; d_{ff} and the number of layers. The number of training tuples can be set by early-stopping or a time budget; §7.4’s results suggest starting with a few to 10+ million.

7.6 Update Strategies

NeuroCard handles new data by either retraining, or taking additional gradient steps, i.e., incremental training. To test both strategies, we simulate the practice of *time-ordered partition appends*: table title is range-partitioned on a year column into 5 partitions. Each partition defines a distinct snapshot of the entire database and the full join, so running the same set of queries at different partition count yields 5 sets of true cardinalities. We compare three update strategies, all of which are trained fully for 7M tuples after the first ingest: (1) *stale*, trained once on the first snapshot and never updated, (2) *fast update*, incrementally updated after each new ingest on 1% of original samples (70K), and (3) *retrain*, using 100% of original samples (7M) after each ingest. We also show the latency required to perform additional gradient steps.

Results are shown in Figure 8. Without update, the stale NeuroCard significantly degrades in accuracy, which is expected as each partition adds a significant amount of new information. A fast updated NeuroCard recovers most of the accuracy, incurring a

minimal overhead. Even fully retraining only requires a few minutes and yields the highest accuracy. Both the statistical efficiency (number of tuples needed vs. accuracy) and the physical efficiency of NeuroCard contribute to these highly practical update strategies.

8 RELATED WORK

Unsupervised data-driven cardinality estimators. This family approximates the data distribution and dates back to System R’s use of 1D histograms [34]. The quality of the density model used has seen steady improvements throughout the years:

Classical methods. Multidimensional histograms [10, 25, 30, 31] are more precise than 1D histograms by capturing inter-column correlations. Starting from early 2000s, graphical models were proposed for either single-table or join cardinality estimation [3, 8, 40]. These density models tradeoff precision for efficiency by assuming conditional or partial independence, and require expensive structure learning (finding the best model structure given a dataset).

Sum-product networks. SPNs, a tree-structured density estimator, were proposed about 10 years ago [29]. Each leaf is a coarse histogram of a slice of an attribute, and each intermediate layer uses either \times and $+$ to combine children information. Due to their heuristics (e.g., inter-slice independence), SPNs have *limited expressiveness*: there exists simple distributions that cannot be efficiently captured by SPNs of any depth [24]. DeepDB [12] is a recent cardinality estimator that uses SPNs. NeuroCard is similar to DeepDB in the following aspects. (S1) Both works use the formulation of learning the full outer join of multiple tables. (S2) Our “schema subsetting” capability builds on their querying algorithms.

NeuroCard differs from DeepDB in the following. (D1) *Modern density model*: NeuroCard’s choice of a deep autoregressive model is a universal function approximator hence fundamentally more expressive. Unlike SPNs, no independence assumption is made in the modeling. (D2) *Correlations learned*: NeuroCard argues for capturing as much correlation as possible across tables, and proposes learning the full outer join of all tables of a schema. DeepDB, due to limited expressiveness, learns multiple SPNs, each on a table subset ($\sim 1\text{--}3$ tables) chosen by correlation tests. Conditional independence is assumed across table subsets. (D3) *Correct sampling*: NeuroCard identifies the key requirement of sampling from the data distribution of joins in an unbiased fashion. In contrast, DeepDB obtains join tuples either from full computation or IBJS which samples from a biased distribution. Due to these differences, NeuroCard outperforms DeepDB by up to $34\times$ in accuracy (§7).

Deep autoregressive models. A breakthrough in density estimation, deep AR models are the current state-of-the-art density models from the ML community [4, 6, 32, 43]. They tractably learn complex, high-dimensional distributions in a neural net, capturing all possible correlations among attributes. Distinctively, AR models provide access to all conditional distributions among input attributes. Naru [48] is a single-table cardinality estimator that uses a deep AR model. By accessing conditional distributions, Naru proposes efficient algorithms to integrate over an AR model, thereby producing selectivity estimates. NeuroCard builds on single-table Naru and overcomes the unique challenges (§2) to support joins.

Supervised query-driven cardinality estimators. Leveraging past or collected queries to improve estimates dates back to LEO [37]. Interest in this approach has seen a resurgence partly due to an abundance of query logs [44] or better function approximators (neural networks) [15, 38] that map featurized queries to predicted cardinalities. Hybrid methods that leverage query feedback to improve density modeling have also been explored, e.g., KDE [11, 14] and mixture of uniforms [49]. Supervised estimators can easily leverage query feedback, handle complex predicates (e.g., UDFs), and are usually more lightweight [5]. NeuroCard has demonstrated superior estimation accuracy to representatives in this family, while being fundamentally more robust since it is not affected by out-of-distribution queries. Complex predicates can also be handled by executing on tuples sampled from NeuroCard’s learned distribution.

Join sampling. Extensive research has studied join sampling, a fundamental problem in databases. NeuroCard leverages a state-of-the-art join sampler to obtain training tuples representative of a join. NeuroCard adopts the linear-time Exact Weight algorithm from Zhao *et al.* [50], which is among the top-performing samplers they study. This algorithm provides uniform and independent samples, just as NeuroCard requires. NeuroCard may further leverage their extensions to support cyclic join schemas. While IBJS [20] and Wander Join [22] provide unbiased estimators for counts and aggregates, they do not provide uniform samples of a join and thus are unsuitable for collecting training data. Lastly, we show that it is advantageous to layer a modern density model on join samples.

Learned database components. A great deal of work has recently applied either classical ML or modern deep learning to various database components, e.g., indexing [17], data layout [47], and query optimization [18, 23, 39]. NeuroCard can be seen as a versatile *core* that can benefit any query engine, learned or not learned. Being able to model inter-table and inter-column correlations without any independence assumptions, NeuroCard’s use may go beyond query optimization to other tasks that require an *understanding of tables and attributes* (e.g., data imputation [45] or indexing [46]).

9 CONCLUSION

NeuroCard is built on a simple idea: learn the correlations across all tables in a database without making any independence assumptions. NeuroCard applies established techniques from join sampling and deep self-supervised learning to cardinality estimation, a fundamental problem in query optimization. It learns from data—just like classical data-driven estimators—but captures all possible inter-table correlations in a probabilistic model: p_{θ} (all tables). To our knowledge, NeuroCard is the first cardinality estimator to achieve assumption-free probabilistic modeling of more than a dozen tables. NeuroCard achieves state-of-the-art accuracy for join cardinality estimation ($4\text{--}34\times$ better than prior methods) using a single per-schema model that is both compact and efficient to learn.

ACKNOWLEDGMENTS

We thank Joe Hellerstein for fruitful discussions and guidance, and Michael Whittaker, Richard Liaw, and Chenggang Wu for their insightful comments on this paper.

REFERENCES

- [1] Michael Armbrust, Reynold S. Xin, Cheng Lian, Yin Huai, Davies Liu, Joseph K. Bradley, Xiangrui Meng, Tomer Kaftan, Michael J. Franklin, Ali Ghodsi, and Matei Zaharia. 2015. Spark SQL: Relational Data Processing in Spark. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data* (Melbourne, Victoria, Australia) (SIGMOD '15). ACM, New York, NY, USA, 1383–1394.
- [2] Kaushik Chakrabarti, Minos Garofalakis, Rajeev Rastogi, and Kyuseok Shim. 2001. Approximate query processing using wavelets. *The VLDB Journal* 10, 2-3 (2001), 199–223.
- [3] Amol Deshpande, Minos Garofalakis, and Rajeev Rastogi. 2001. Independence is good: Dependency-based histogram synopses for high-dimensional data. *ACM SIGMOD Record* 30, 2 (2001), 199–210.
- [4] Conor Durkan and Charlie Nash. 2019. Autoregressive Energy Machines. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.), Vol. 97. PMLR, Long Beach, California, USA, 1735–1744.
- [5] Anshuman Dutt, Chi Wang, Azade Nazi, Srikanth Kandula, Vivek Narasayya, and Surajit Chaudhuri. 2019. Selectivity estimation for range predicates using lightweight models. *Proceedings of the VLDB Endowment* 12, 9 (2019), 1044–1057.
- [6] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. 2015. MADE: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*. 881–889.
- [7] Lise Getoor, Nir Friedman, Daphne Koller, and Benjamin Taskar. 2001. Learning probabilistic models of relational structure. In *ICML*, Vol. 1. 170–177.
- [8] Lise Getoor, Benjamin Taskar, and Daphne Koller. 2001. Selectivity estimation using probabilistic models. In *ACM SIGMOD Record*, Vol. 30. ACM, 461–472.
- [9] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [10] Dimitrios Gunopulos, George Kollios, Vassilis J Tsotras, and Carlotta Domeniconi. 2005. Selectivity estimators for multidimensional range queries over real attributes. *The VLDB Journal* 14, 2 (2005), 137–154.
- [11] Max Heimel, Martin Kiefer, and Volker Markl. 2015. Self-Tuning, GPU-Accelerated Kernel Density Models for Multidimensional Selectivity Estimation. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data* (SIGMOD '15). ACM, New York, NY, USA, 1477–1492.
- [12] Benjamin Hilprecht, Andreas Schmidt, Moritz Kulessa, Alejandro Molina, Kristian Kersting, and Carsten Binnig. 2020. DeepDB: Learn from Data, not from Queries! *Proceedings of the VLDB Endowment* 13, 7 (2020), 992–1005.
- [13] Hilprecht et al. 2020. Github repository, deepdb-public. github.com/DataManagementLab/deepdb-public. [Online; accessed April, 2020].
- [14] Martin Kiefer, Max Heimel, Sebastian Breß, and Volker Markl. 2017. Estimating join selectivities using bandwidth-optimized kernel density models. *Proceedings of the VLDB Endowment* 10, 13 (2017), 2085–2096.
- [15] 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 2019, 9th Biennial Conference on Innovative Data Systems Research, Asilomar, CA, USA, January 13-16, 2019, Online Proceedings*.
- [16] Kipf et al. 2019. Github repository, learnedcardinalities. github.com/andreaskipf/learnedcardinalities. [Online; accessed April, 2020].
- [17] Tim Kraska, Alex Beutel, Ed H Chi, Jeffrey Dean, and Neoklis Polyzotis. 2018. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*. ACM, 489–504.
- [18] Sanjay Krishnan, Zongheng Yang, Ken Goldberg, Joseph Hellerstein, and Ion Stoica. 2018. Learning to optimize join queries with deep reinforcement learning. *arXiv preprint arXiv:1808.03196* (2018).
- [19] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter Boncz, Alfons Kemper, and Thomas Neumann. 2015. How good are query optimizers, really? *Proceedings of the VLDB Endowment* 9, 3 (2015), 204–215.
- [20] Viktor Leis, Bernhard Radke, Andrey Gubichev, Alfons Kemper, and Thomas Neumann. 2017. Cardinality Estimation Done Right: Index-Based Join Sampling.. In *CIDR*.
- [21] Viktor Leis, Bernhard Radke, Andrey Gubichev, Atanas Mirchev, Peter Boncz, Alfons Kemper, and Thomas Neumann. 2018. Query optimization through the looking glass, and what we found running the join order benchmark. *The VLDB Journal* (2018), 1–26.
- [22] Feifei Li, Bin Wu, Ke Yi, and Zhuoyue Zhao. 2016. Wander join: Online aggregation via random walks. In *Proceedings of the 2016 International Conference on Management of Data*. 615–629.
- [23] Ryan Marcus, Parimarjan Negi, Hongzi Mao, Chi Zhang, Mohammad Alizadeh, Tim Kraska, Olga Papaemmanouil, and Nesime Tatbul. 2019. Neo: A Learned Query Optimizer. *PVLDB* 12, 11 (2019), 1705–1718.
- [24] James Martens and Venkatesh Medabalimi. 2014. On the expressive efficiency of sum product networks. *arXiv preprint arXiv:1411.7717* (2014).
- [25] M Muralikrishna and David J DeWitt. 1988. Equi-depth multidimensional histograms. In *ACM SIGMOD Record*, Vol. 17. ACM, 28–36.
- [26] Kevin P Murphy. 2012. *Machine learning: a probabilistic perspective*. MIT press.
- [27] Neural Relation Understanding (Naru). 2020. Github repository, [naru. github.com/naru-project/naru](https://github.com/naru-project/naru). [Online; accessed April, 2020].
- [28] Patrick O’Neil and Dallan Quass. 1997. Improved query performance with variant indexes. In *Proceedings of the 1997 ACM SIGMOD international conference on Management of data*. 38–49.
- [29] Hoifung Poon and Pedro Domingos. 2011. Sum-product networks: A new deep architecture. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*. IEEE, 689–690.
- [30] Viswanath Poosala, Peter J. Haas, Yannis E. Ioannidis, and Eugene J. Shekita. 1996. Improved Histograms for Selectivity Estimation of Range Predicates. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data* (Montreal, Quebec, Canada) (SIGMOD '96). ACM, New York, NY, USA, 294–305.
- [31] Viswanath Poosala and Yannis E Ioannidis. 1997. Selectivity estimation without the attribute value independence assumption. In *VLDB*, Vol. 97. 486–495.
- [32] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. [URL https://openai.com/blog/better-language-models](https://openai.com/blog/better-language-models) (2019).
- [33] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P. Kingma. 2017. PixelCNN++: Improving the PixelCNN with Discretized Logistic Mixture Likelihood and Other Modifications. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- [34] P Griffiths 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 *Proceedings of the 1979 ACM SIGMOD international conference on Management of data*. ACM, 23–34.
- [35] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 1715–1725.
- [36] R. Sethi, M. Traverso, D. Sundstrom, D. Phillips, W. Xie, Y. Sun, N. Yegitbasi, H. Jin, E. Hwang, N. Shingte, and C. Berner. 2019. Presto: SQL on Everything. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. 1802–1813.
- [37] Michael Stillger, Guy M Lohman, Volker Markl, and Mokhtar Kandil. 2001. LEO-DB2’s learning optimizer. In *VLDB*, Vol. 1. 19–28.
- [38] Ji Sun and Guoliang Li. 2019. An end-to-end learning-based cost estimator. *Proceedings of the VLDB Endowment* 13, 3 (2019), 307–319.
- [39] Immanuel Trummer, Junxiang Wang, Deepak Maram, Samuel Moseley, Saehan Jo, and Joseph Antonakakis. 2019. SkinnerDB: Regret-Bounded Query Evaluation via Reinforcement Learning. In *Proceedings of the 2019 International Conference on Management of Data (SIGMOD '19)*. ACM, New York, NY, USA, 1153–1170.
- [40] Kostas Tzoumas, Amol Deshpande, and Christian S Jensen. 2011. Lightweight graphical models for selectivity estimation without independence assumptions. *Proceedings of the VLDB Endowment* 4, 11 (2011), 852–863.
- [41] Kostas Tzoumas, Amol Deshpande, and Christian S Jensen. 2013. Efficiently adapting graphical models for selectivity estimation. *The VLDB Journal* 22, 1 (2013), 3–27.
- [42] Aaron Van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. WaveNet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499* (2016).
- [43] 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 *Advances in neural information processing systems*. 5998–6008.
- [44] Chenggang Wu, Alekh Jindal, Saeed Amizadeh, Hiren Patel, Wangchao Le, Shi Qiao, and Sriram Rao. 2018. Towards a learning optimizer for shared clouds. *Proceedings of the VLDB Endowment* 12, 3 (2018), 210–222.
- [45] Richard Wu, Aoqian Zhang, Ihab Ilyas, and Theodoros Rekatsinas. 2020. Attention-based Learning for Missing Data Imputation in HoloClean. *Proceedings of Machine Learning and Systems* (2020), 307–325.
- [46] Yingjun Wu, Jia Yu, Yuanyuan Tian, Richard Sidle, and Ronald Barber. 2019. Designing succinct secondary indexing mechanism by exploiting column correlations. In *Proceedings of the 2019 International Conference on Management of Data*. 1223–1240.
- [47] Zongheng Yang, Badrish Chandramouli, Chi Wang, Johannes Gehrke, Yanan Li, Umar Farooq Minhas, Per-Åke Larson, Donald Kossmann, and Rajeev Acharya. 2020. Qd-tree: Learning Data Layouts for Big Data Analytics. In *Proceedings of the 2020 International Conference on Management of Data (SIGMOD '20)*.
- [48] 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. *Proceedings of the VLDB Endowment* 13, 3 (2019), 279–292.
- [49] Barzan Mozafari Yongjoo Park, Shucheng Zhong. 2020. QuickSel: Quick Selectivity Learning with Mixture Models. *SIGMOD* (2020).
- [50] Zhuoyue Zhao, Robert Christensen, Feifei Li, Xiao Hu, and Ke Yi. 2018. Random sampling over joins revisited. In *Proceedings of the 2018 International Conference on Management of Data*. 1525–1539.