

QuickSel: Quick Selectivity Learning with Mixture Models

Yongjoo Park^{1*}, Shucheng Zhong^{2*}, Barzan Mozafari²
University of Illinois at Urbana-Champaign¹ University of Michigan²
yongjoo@illinois.edu {joezhong,mozafari}@umich.edu

ABSTRACT

Estimating the selectivity of a query is a key step in almost any cost-based query optimizer. Most of today’s databases rely on histograms or samples that are periodically refreshed by re-scanning the data as the underlying data changes. Since frequent scans are costly, these statistics are often stale and lead to poor selectivity estimates. As an alternative to scans, *query-driven histograms* have been proposed, which refine the histograms based on the actual selectivities of the observed queries. Unfortunately, these approaches are either too costly to use in practice—i.e., require an exponential number of buckets—or quickly lose their advantage as they observe more queries.

In this paper, we propose a *selectivity learning* framework, called QuickSel, which falls into the query-driven paradigm but does not use histograms. Instead, it builds an internal *model* of the underlying data, which can be refined significantly faster (e.g., only 1.9 milliseconds for 300 queries). This fast refinement allows QuickSel to continuously learn from *each query* and yield increasingly more accurate selectivity estimates over time. Unlike query-driven histograms, QuickSel relies on a mixture model and a new optimization algorithm for training its model. Our extensive experiments on two real-world datasets confirm that, given the same target accuracy, QuickSel is 34.0×–179.4× faster than state-of-the-art query-driven histograms, including ISOMER and STHoles. Further, given the same space budget, QuickSel is 26.8%–91.8% more accurate than periodically-updated histograms and samples, respectively.

*These authors contributed equally to this work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGMOD’20, June 14–19, 2020, Portland, OR, USA
© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6735-6/20/06...\$15.00
<https://doi.org/10.1145/3318464.3389727>

ACM Reference Format:

Yongjoo Park, Shucheng Zhong, Barzan Mozafari. 2020. QuickSel: Quick Selectivity Learning with Mixture Models. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD’20)*, June 14–19, 2020, Portland, OR, USA. ACM, New York, NY, USA, Article 4, 17 pages. <https://doi.org/10.1145/3318464.3389727>

1 INTRODUCTION

Estimating the *selectivity* of a query—the fraction of input tuples that satisfy the query’s predicate—is a fundamental component in cost-based query optimization, including both traditional RDBMSs [2, 3, 7, 9, 83] and modern SQL-on-Hadoop engines [41, 88]. The estimated selectivities allow the query optimizer to choose the cheapest access path or query plan [53, 90].

Today’s databases typically rely on histograms [2, 7, 9] or samples [83] for their selectivity estimation. These structures need to be populated in advance by performing costly table scans. However, as the underlying data changes, they quickly become stale and highly inaccurate. This is why they need to be updated periodically, creating additional costly operations in the database engine (e.g., `ANALYZE table`).¹

To address the shortcoming of scan-based approaches, numerous proposals for query-driven histograms have been introduced, which continuously correct and refine the histograms based on the actual selectivities observed after running each query [11, 12, 19, 52, 66, 75, 86, 93, 96]. There are two approaches to query-driven histograms. The first approach [11, 12, 19, 66], which we call *error-feedback histograms*, recursively splits existing buckets (both boundaries and frequencies) for every distinct query observed, such that their error is minimized for the latest query. Since the error-feedback histograms do not minimize the (average) error across multiple queries, their estimates tend to be much less accurate.

To achieve a higher accuracy, the second approach is to compute the bucket frequencies based on the maximum entropy principle [52, 75, 86, 93]. However, this approach

¹Some database systems [9] automatically update their statistics when the number of modified tuples exceeds a threshold.

Table 1: The differences between query-driven histograms [52, 74, 75, 86, 93] and our method (QuickSel)

	Query-driven Histograms	QuickSel (ours)	Our Contribution
Model	histograms (non-overlapping buckets)	mixture models (overlapping subpopulations)	Employs a new expressive model → no exponential growth of complexity
Training	<i>maximum entropy</i> solved by <i>iterative scaling</i>	<i>min difference from a uniform distribution</i> solved <i>analytically</i>	A new optimization objective and its reduction to quadratic programming (solved analytically) → fast training and model refinements

(which is also the state-of-the-art) requires solving an optimization problem, which quickly becomes prohibitive as the number of observed queries (and hence, number of buckets) grows. Unfortunately, one cannot simply prune the buckets in this approach, as it will break the underlying assumptions of their optimization algorithm (called *iterative scaling*, see §2.3 for details). Therefore, they prune the *observed queries* instead in order to keep the optimization overhead feasible in practice. However, this also means discarding data that could be used for learning a more accurate distribution.

Our Goal We aim to develop a new framework for selectivity estimation that can quickly refine its *model* after observing each query, thereby producing increasingly more accurate estimates over time. We call this new framework *selectivity learning*. We particularly focus on designing a low-overhead method that can scale to a large number of observed queries without requiring an exponential number of buckets.

Our Model To overcome the limitations of query-driven histograms, we use a *mixture model* [17] to capture the unknown distribution of the data. A mixture model is a probabilistic model to approximate an arbitrary probability density function (pdf), say $f(x)$, using a combination of simpler pdfs:

$$f(x) = \sum_{z=1}^m h(z) g_z(x) \quad (1)$$

where $g_z(x)$ is the z -th simpler pdf and $h(z)$ is its corresponding weight. The subset of the data that follows each of the simpler pdfs is called a *subpopulation*. Since the subpopulations are allowed to overlap with one another, a mixture model is strictly more expressive than histograms. In fact, it is shown that mixture models can achieve a higher accuracy than histograms [24], which is confirmed by our empirical study (§5.5). To the best of our knowledge, we are the first to propose a mixture model for selectivity estimation.²

Challenges Using a mixture model for selectivity learning requires finding optimal parameter values for $h(z)$ and

$g_z(x)$; however, this optimization (a.k.a. training) is challenging for two reasons.

First, the training process aims to find parameters that maximize the model *quality*, defined as $\int Q(f(x)) dx$ for some metric of quality Q (e.g., entropy). However, computing this integral is non-trivial for a mixture model since its subpopulations may overlap in arbitrary ways. That is, the combinations of m subpopulations can create 2^m distinct ranges, each with a potentially different value of $f(x)$. As a result, naively computing the quality of a mixture model quickly becomes intractable as the number of observed queries grows.

Second, the outer optimization algorithms are often iterative (e.g., iterative scaling, gradient descent), which means they have to repeatedly evaluate the model quality as they search for optimal parameter values. Thus, even when the model quality can be evaluated relatively efficiently, the overall training/optimization process can be quite costly.

Our Approach First, to ensure the efficiency of the model quality evaluation, we propose a new optimization objective. Specifically, we find the parameter values that minimize the L_2 distance (or equivalently, *mean squared error*) between the mixture model and a uniform distribution, rather than maximizing the entropy of the model (as pursued by previous work [52, 74, 75, 86, 93]). As described above, directly computing the quality of a mixture model involves costly integrations over 2^m distinct ranges. However, when minimizing the L_2 distance, the 2^m integrals can be reduced to only m^2 multiplications, hence greatly reducing the complexity of the model quality evaluation. Although minimizing the L_2 distance is much more efficient than maximizing the entropy, these two objectives are closely related (see our report [82] for a discussion).

In addition, we adopt a non-conventional variant of mixture models, called a *uniform mixture model*. While uniform mixture models have been previously explored in limited settings (with only a few subpopulations) [26, 36], we find that they are quite appropriate in our context as they allow for efficient computations of the L_2 distance. That is, with this choice, we can evaluate the quality of a model by only using min, max, and multiplication operations (§3.2). Finally, our optimization can be expressed as a standard *quadratic program*, which still requires an iterative procedure.

²Our mixture model also differs from *kernel density estimation* [18, 35, 40], which scans the actual data, rather than analyzing observed queries.

Therefore, to avoid the costly iterative optimization, we also devise an analytic solution that can be computed more efficiently. Specifically, in addition to the standard reduction (i.e., moving some of the original constraints to the objective clause as penalty terms), we completely relax the positivity constraints for $f(x)$, exploiting the fact that they will be naturally satisfied in the process of approximating the true distribution of the data. With these modifications, we can solve for the solution analytically by setting the gradient of the objective function to zero. This simple transformation speeds up the training by $1.5\times$ – $17.2\times$. In addition, since our analytic solution requires a constant number of operations, the training time is also consistent across different datasets and workloads.

Using these ideas, we have developed a first prototype of our selectivity learning proposal, called QuickSel, which allows for extremely fast model refinements. As summarized in Table 1, QuickSel differs from—and considerably improves upon—query-driven histograms [11, 52, 66, 75, 86, 93] in terms of both modeling and training (see §7 for a detailed comparison).

Contributions We make the following contributions:

1. We propose the first mixture model-based approach to selectivity estimation (§3).
2. For training the mixture model, we design a constrained optimization problem (§4.1).
3. We show that the proposed optimization problem can be reduced (from an exponentially complex one) to a quadratic program and present further optimization strategies for solving it (§4.2).
4. We conduct extensive experiments on two real-world datasets to compare QuickSel’s performance and state-of-the-art selectivity estimation techniques (§5).

2 PRELIMINARIES

In this section, we first define relevant notations in §2.1 and then formally define the problem of *query-driven selectivity estimation* in §2.2. Next, in §2.3, we discuss the drawbacks of previous approaches.

2.1 Notations

Table 2 summarizes the notations used throughout this paper.

Set Notations T is a relation that consists of d real-valued columns C_1, \dots, C_d .³ The range of values in C_i is $[l_i, u_i]$ and the cardinality (i.e., row count) of T is $N=|T|$. The tuples in T are denoted by x_1, \dots, x_N , where each x_i is a size- d vector that belongs to $B_0 = [l_1, u_1] \times \dots \times [l_d, u_d]$. Geometrically, B_0 is the area bounded by a hyperrectangle whose bottom-left corner is (l_1, \dots, l_d) and top-right corner is (u_1, \dots, u_d) . The

³Handling integer and categorical columns is discussed in §2.2.

Table 2: Notations.

Symbol	Meaning
T	a table (or a relation)
C_i	i -th column (or an attribute) of T ; $i = 1, \dots, d$
$ T $	the number of tuples in T
$[l_i, u_i]$	the range of the values in C_i
x	a tuple of T
B_0	the domain of x ; $[l_1, u_1] \times \dots \times [l_d, u_d]$
P_i	i -th predicate
B_i	hyperrectangle range for the i -th predicate
$ B_i $	the size (of the area) of B_i
$x \in B_i$	x belongs to B_i ; thus, satisfies P_i
$I(\cdot)$	indicator function that returns 1 if its argument is true and 0 otherwise
s_i	the selectivity of P_i for T
(P_i, s_i)	i -th observed query
n	the total number of observed queries
$f(x)$	probability density function of tuple x (of T)

size of B_0 can thus be computed as $|B_0|=(u_1 - l_1) \times \dots \times (u_d - l_d)$.

Predicates We use P_i to denote the (selection) predicate of the i -th query on T . In this paper, a predicate is a conjunction⁴ of one or more *constraints*. Each constraint is a *range constraint*, which can be one-sided (e.g., $3 \leq C_1$) or two-sided (e.g., $-3 \leq C_1 \leq 10$). This range can be extended to also handle equality constraints on categorical data (see §2.2). Each predicate P_i is represented by a hyperrectangle B_i . For example, a constraint “ $1 \leq C_1 \leq 3$ AND $2 \leq C_2$ ” is represented by a hyperrectangle $(1, 3) \times (2, u_2)$, where u_2 is the upper-bound of C_2 . We use P_o to denote an empty predicate, i.e., one that selects all tuples.

Selectivity The *selectivity* s_i of P_i is defined as the *fraction* of the rows of T that satisfy the predicate. That is, $s_i = (1/N) \sum_{k=1}^N I(x_k \in B_i)$, where $I(\cdot)$ is the indicator function. A pair (P_i, s_i) is referred to as an *observed query*.⁵ Without loss of generality, we assume that n queries have been observed for T and seek to estimate s_{n+1} . Finally, we use $f(x)$ to denote the joint probability density function of tuple x (that has generated tuples of T).

2.2 Problem Statement

Next, we formally state the problem:

Problem 1 (Query-driven Selectivity Estimation) Consider a set of n observed queries $(P_1, s_1), \dots, (P_n, s_n)$ for T . By definition, we have the following for each $i = 1, \dots, n$:

$$\int_{x \in B_i} f(x) dx = s_i$$

⁴See §2.2 for a discussion of disjunctions and negations.

⁵This pair is also referred to as an *assertion* by prior work [85].

Then, our goal is to build a model of $f(x)$ that can estimate the selectivity s_{n+1} of a new predicate P_{n+1} .

Initially, before any query is observed, we can conceptually consider a default query $(P_0, 1)$, where all tuples are selected and hence, the selectivity is 1 (i.e., no predicates).

Discrete and Categorical Values Problem 1 can be extended to support discrete attributes (e.g., integers, characters, categorical values) and equality constraints on them, as follows. Without loss of generality, suppose that C_i contains the integers in $\{1, 2, \dots, b_i\}$. To apply the solution to Problem 1, it suffices to (conceptually) treat these integers as real values in $[1, b_i + 1]$ and then convert the original constraints on the integer values into range constraints, as follows. A constraint of the form “ $C_i = k$ ” will be converted to a range constraint of the form $k \leq C_i < k + 1$. Mathematically, this is equivalent to replacing a probability mass function with a probability density function defined using *dirac delta functions*.⁶ Then, the summation of the original probability mass function can be converted to the integration of the probability density function. String data types (e.g., char, varchar) and their equality constraints can be similarly supported, by conceptually mapping each string into an integer (preserving their order) and applying the conversion described above for the integer data type.

Supported Queries Similar to prior work [11, 19, 52, 66, 74, 75, 86, 93], we support selectivity estimation for predicates with conjunctions, negations, and disjunctions of range and equality constraints on numeric and categorical columns. We currently do not support wildcard constraints (e.g., LIKE ‘*word*’), EXISTS constraints, or ANY constraints. In practice, often a *fixed selectivity* is used for unsupported predicates, e.g., 3.125% in Oracle [83].

To simplify our presentation, we focus on conjunctive predicates. However, negations and disjunctions can also be easily supported. This is because our algorithm only requires the ability to compute the intersection size of pairs of predicates P_i and P_j , which can be done by converting $P_i \wedge P_j$ into a disjunctive normal form and then using the inclusion-exclusion principle to compute its size.

As in the previous work, we focus our presentation on predicates on a single relation. However, any selectivity estimation technique for a single relation can be applied to estimating selectivity of a join query whenever the predicates on the individual relations are independent of the join conditions [7, 41, 90, 98].

⁶A dirac delta function $\delta(x)$ outputs ∞ if $x = 0$ and outputs 0 otherwise while satisfying $\int \delta(x) dx = 1$.




Figure 1: Bucket creation for query-driven histograms. P_3 is the range of a new predicate. The existing buckets (for P_1 and P_2) are split into multiple buckets. The total number of buckets may grow *exponentially* as more queries are observed.

2.3 Why not Query-driven Histograms

In this section, we briefly describe how query-driven histograms work [11, 19, 52, 66, 75, 86, 93], and then discuss their limitations, which motivate our work.

How Query-driven Histograms Work To approximate $f(x)$ (defined in Problem 1), query-driven histograms adjust their bucket boundaries and bucket frequencies according to the queries they observe. Specifically, they first determine bucket boundaries (*bucket creation* step), and then compute their frequencies (*training* step), as described next.

1. *Bucket Creation:* Query-driven histograms determine their bucket boundaries based on the given predicate’s ranges [11, 52, 75, 86, 93]. If the range of a later predicate overlaps with that of an earlier predicate, they split the bucket(s) created for the earlier one into two or more buckets in order to ensure that the buckets do not overlap with one another. Figure 1 shows an example of this bucket splitting operation.
2. *Training:* After creating the buckets, query-driven histograms assign frequencies to those buckets. Early work [11, 66] determines bucket frequencies in the process of bucket creations. That is, when a bucket is split into two or more, the frequency of the original bucket is also split (or adjusted), such that it minimizes the estimation error for the latest observed query.

However, since this process does not minimize the (average) error across multiple queries, their estimates are much less accurate. More recent work [52, 75, 86, 93] has addressed this limitation by explicitly solving an optimization problem based on the maximum entropy principle. That is, they search for bucket frequencies that maximize the *entropy* of the distribution while remaining consistent with the actual selectivities observed.

Although using the maximum entropy principle will lead to highly accurate estimates, it still suffers from two key limitations.

Limitation 1: Exponential Number of Buckets Since existing buckets may split into multiple ones for each new observed query, the number of buckets can potentially grow

exponentially as the number of observed queries grows. For example, in our experiment in §5.5, the number of buckets was 22,370 for 100 observed queries, and 318,936 for 300 observed queries. Unfortunately, the number of buckets directly affects the training time. Specifically, using *iterative scaling*—the optimization algorithm used by all previous work [52, 74, 75, 86, 93]—the cost of each iteration grows linearly with the number of variables (i.e., the number of buckets). This means that the cost of each iteration can grow exponentially with the number of observed queries.

As stated in §1, we address this problem by employing a *mixture model*, which can express a probability distribution more effectively than query-driven histograms. Specifically, our empirical study in §5.5 shows that—using the same number of parameters—a mixture model achieves considerably more accurate estimates than histograms.

Limitation 2: Non-trivial Bucket Merge/Pruning Given that query-driven histograms [75, 93] quickly become infeasible due to their large number of buckets, one might consider merging or pruning the buckets in an effort to reduce their training times. However, merging or pruning the histogram buckets violates the assumption used by their optimization algorithms, i.e., iterative scaling. Specifically, iterative scaling relies on the fact that a bucket is either *completely included* in a query’s predicate range or *completely outside* of it.⁷ That is, no partial overlap is allowed. This property must hold for each of the n predicates. However, merging some of the buckets will inevitably cause partial overlaps (between predicate and histogram buckets). For interested readers, we have included a more detailed explanation of why iterative scaling requires this assumption in our technical report [82].

3 QUICKSEL: MODEL

This section presents how QuickSel models the population distribution and estimates the selectivity of a new query. QuickSel’s model relies on a probabilistic model called a *mixture model*. In §3.1, we describe the mixture model employed by QuickSel. §3.2 describes how to estimate the selectivity of a query using the mixture model. §3.3 describes the details of QuickSel’s mixture model construction.

3.1 Uniform Mixture Model

A mixture model is a probabilistic model that expresses a (complex) probability density function (of the population) as a combination of (simpler) probability density functions (of *subpopulations*). The population distribution is the one that

⁷For example, this property is required for the transition from Equation (6) to Equation (7) in [75].

generates the tuple x of T . The subpopulations are internally managed by QuickSel to best approximate $f(x)$.

Uniform Mixture Model QuickSel uses a type of mixture model, called *the uniform mixture model*. The uniform mixture model represents a population distribution $f(x)$ as a weighted summation of multiple uniform distributions, $g_z(x)$ for $z = 1, \dots, m$. Specifically,

$$f(x) = \sum_{z=1}^m h(z) g_z(x) = \sum_{z=1}^m w_z g_z(x) \quad (2)$$

where $h(z)$ is a categorical distribution that determines the weight of the z -th subpopulation, and $g_z(x)$ is the probability density function (which is a uniform distribution) for the z -th subpopulation. The support of $h(z)$ is the integers ranging from 1 to m ; $h(z) = w_z$. The support for $g_z(x)$ is represented by a hyperrectangle G_z . Since $g_z(x)$ is a uniform distribution, $g_z(x) = 1/|G_z|$ if $x \in G_z$ and 0 otherwise. The locations of G_z and the values of w_z are determined in the training stage (§4). In the remainder of this section (§3), we assume that G_z and w_z are given.

Benefit of Uniform Mixture Model The uniform mixture model was studied early in the statistics community [26, 36]; however, recently, a more complex model called *the Gaussian mixture model* has received more attention [17, 84, 110].⁸ The Gaussian mixture model uses a Gaussian distribution for each subpopulation; the smoothness of its probability density function (thus, differentiable) makes the model more appealing when gradients need to be computed. Nevertheless, we *intentionally* use the uniform mixture model for QuickSel due to its computational benefit in the training process, as we describe below.

As will be presented in §4.2, QuickSel’s training involves the computations of the intersection size between two subpopulations, for which the essential operation is evaluating the following integral:

$$\int g_{z1}(x) g_{z2}(x) dx$$

Evaluating the above expression for multivariate Gaussian distributions, e.g., $g_{z1}(x) = \exp(-x^T \Sigma^{-1} x) / \sqrt{(2\pi)^d |\Sigma|}$, requires numerical approximations [31, 50], which are either slow or inaccurate. In contrast, the intersection size between two hyperrectangles can be exactly computed by simple min, max, and multiplication operations.

3.2 Selectivity Estimation with UMM

For the uniform mixture model, computing the selectivity of a predicate P_i is straightforward:

⁸There are other variants of mixture models [15, 78].

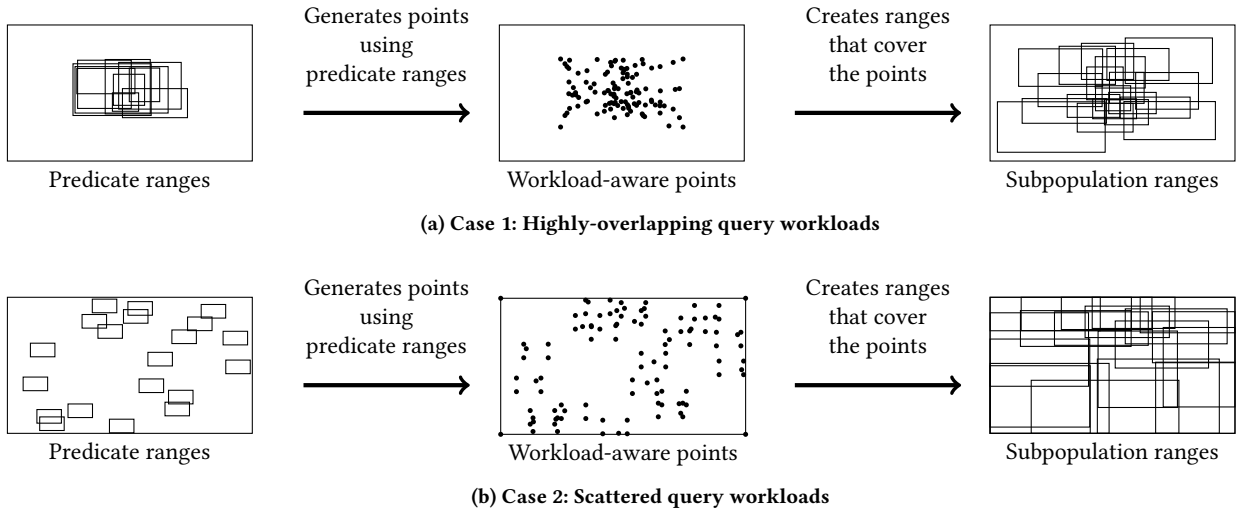


Figure 2: QuickSel’s subpopulation creation. Due to the property of mixture model (i.e., subpopulations may overlap with one another), creating subpopulations is straightforward for diverse query workloads.

$$\begin{aligned} \int_{B_i} f(x) dx &= \int_{B_i} \sum_{z=1}^m w_z g_z(x) dx = \sum_{z=1}^m w_z \int_{B_i} g_z(x) dx \\ &= \sum_{z=1}^m w_z \int \frac{1}{|G_z|} I(x \in G_z \cap B_i) dx = \sum_{z=1}^m w_z \frac{|G_z \cap B_i|}{|G_z|} \end{aligned}$$

Recall that both G_z and B_i are represented by hyperrectangles. Thus, their intersection is also a hyperrectangle, and computing its size is straightforward.

3.3 Subpopulations from Observed Queries

We describe QuickSel’s approach to determining the boundaries of G_z for $z = 1, \dots, m$. Note that determining G_z is orthogonal to the model training process, which we describe in §4; thus, even if one devises an alternative approach to creating G_z , our fast training method is still applicable.

QuickSel creates m hyperrectangular ranges⁹ (for the supports of its subpopulations) in a way that satisfies the following simple criterion: if more predicates involve a point x , use a larger number of subpopulations for x . Unlike query-driven histograms, QuickSel can easily pursue this goal by exploiting the property of a mixture model: the supports of subpopulations may overlap with one another.

In short, QuickSel generates multiple points (using predicates) that represent the query workloads and create hyperrectangles that can sufficiently cover those points. Specifically, we propose two approaches for this: a sampling-based one and a clustering-based one. The sampling-based approach is faster; the clustering-based approach is more accurate. Each of these is described in more detail below.

⁹The number m of subpopulations is set to $\min(4 \cdot n, 4000)$, by default.

Sampling-based This approach performs the following operations for creating G_z for $z = 1, \dots, m$.

1. Within each predicate range, generate multiple random points r . Generating a large number of random points increases the consistency; however, QuickSel limits the number to 10 since having more than 10 points did not improve accuracy in our preliminary study.
2. Use simple random sampling to reduce the number of points to m , which serves as the centers of G_z for $z = 1, \dots, m$.
3. The length of the i -th dimension of G_z is set to twice the average of the distances (in the same i -th dimension) to the 10 nearest-neighbor centers.

Figure 2 illustrates how the subpopulations are created using both (1) highly-overlapping query workloads and (2) scattered query workloads. In both cases, QuickSel generates random points to represent the distribution of query workloads, which is then used to create G_z ($z = 1, \dots, m$), i.e., the supports of subpopulations. This sampling-based approach is faster, but it does not ensure the coverage of all random points r . In contrast, the following clustering-based approach ensures that.

Clustering-based The second approach relies on a clustering algorithm for generating hyperrectangles:

1. Do the same as the sampling-based approach.
2. Cluster r into m groups. (We used K-means++.)
3. For each of m groups, we create the smallest hyperrectangle G_z that covers all the points belonging to the group.

Note that since each r belongs to a cluster and we have created a hyperrectangle that fully covers each cluster, the union

of the hyperrectangles covers all r . Our experiments primarily use the sampling-based approach due to its efficiency, but we also compare them empirically in §5.7.

The following section describes how to assign the weights (i.e., $h(z) = w_z$) of these subpopulations.

4 QUICKSEL: MODEL TRAINING

This section describes how to compute the weights w_z of QuickSel’s subpopulations. For training its model, QuickSel finds the model that maximizes uniformity while being consistent with the observed queries. In §4.1, we formulate an optimization problem based on this criteria. Next, §4.2 presents how to solve the problem efficiently.

4.1 Training as Optimization

This section formulates an optimization problem for QuickSel’s training. Let $g_0(x)$ be the uniform distribution with support B_0 ; that is, $g_0(x) = 1/|B_0|$ if $x \in B_0$ and 0 otherwise. QuickSel aims to find the model $f(x)$, such that the difference between $f(x)$ and $g_0(x)$ is minimized while being consistent with the observed queries.

There are many metrics that can measure the distance between two probability density functions $f(x)$ and $g_0(x)$, such as the earth mover’s distance [89], Kullback-Leibler divergence [62], the mean squared error (MSE), the Hellinger distance, and more. Among them, QuickSel uses MSE (which is equivalent to L_2 distance between two distributions) since it enables the reduction of our originally formulated optimization problem (presented shortly; Problem 2) to a quadratic programming problem, which can be solved efficiently by many off-the-shelf optimization libraries [1, 4, 5, 14]. Also, minimizing MSE between $f(x)$ and $g_0(x)$ is closely related to maximizing the entropy of $f(x)$ [52, 75, 86, 93]. See §6 for the explanation of this relationship.

MSE between $f(x)$ and $g_0(x)$ is defined as follows:

$$\text{MSE}(f(x), g_0(x)) = \int (f(x) - g_0(x))^2 dx$$

Recall that the support for $g_0(x)$ is B_0 . Thus, QuickSel obtains the optimal weights by solving the following problem.

Problem 2 (QuickSel’s Training) *QuickSel obtains the optimal parameter \mathbf{w} for its model by solving:*

$$\arg \min_{\mathbf{w}} \int_{x \in B_0} \left(f(x) - \frac{1}{|B_0|} \right)^2 dx \quad (3)$$

$$\text{such that } \int_{B_i} f(x) dx = s_i \text{ for } i = 1, \dots, n \quad (4)$$

$$f(x) \geq 0 \quad (5)$$

Here, (5) ensures $f(x)$ is a proper probability density function.

4.2 Efficient Optimization

We first describe the challenges in solving Problem 2. Then, we describe how to overcome the challenges.

Challenge Solving Problem 2 in a naïve way is computationally intractable. For example, consider a mixture model consisting of (only) two subpopulations represented by G_1 and G_2 , respectively. Then, $\int_{x \in B_0} (f(x) - g_0(x))^2 dx$ is:

$$\int_{x \in G_1 \cap G_2} \left(\frac{w_1 + w_2}{|G_1 \cap G_2|} - g_0(x) \right)^2 dx + \int_{x \in G_1 \cap \neg G_2} (\dots)^2 dx \\ + \int_{x \in \neg G_1 \cap G_2} (\dots)^2 dx + \int_{x \in \neg G_1 \cap \neg G_2} \left(\frac{0}{|\neg G_1 \cap \neg G_2|} - g_0(x) \right)^2 dx$$

Observe that with this approach, we need four separate integrations only for two subpopulations. In general, the number of integrations is $O(2^m)$, which is $O(2^n)$. Thus, this direct approach is computationally intractable.

Conversion One: Quadratic Programming Problem 2 can be solved efficiently by exploiting the property of the distance metric of our choice (i.e., MSE) and the fact that we use uniform distributions for subpopulations (i.e., UMM). The following theorem presents the efficient approach.

Theorem 1 *The optimization problem in Problem 2 can be solved by the following quadratic optimization:*

$$\arg \min_{\mathbf{w}} \mathbf{w}^T Q \mathbf{w} \quad \text{such that } A \mathbf{w} = \mathbf{s}, \quad \mathbf{w} \succeq \mathbf{0}$$

where $(Q)_{ij} = |G_i \cap G_j| / (|G_i| |G_j|)$, and $(A)_{ij} = |B_i \cap G_j| / |G_j|$. The bendy inequality sign (\succeq) means that every element of the vector on the left-hand side is equal to or larger than the corresponding element of the vector on the right-hand side.

PROOF. This theorem can be shown by substituting the definition of QuickSel’s model (Equation (2)) into the probability density function $f(x)$ in Equation (3). Note that minimizing $(f(x) - g_0(x))^2$ is equivalent to minimizing $f(x) (f(x) - 2g_0(x))$, which is also equivalent to minimizing $(f(x))^2$ since $g_0(x)$ is constant over B_0 and $\int f(x) dx = 1$.

The integration of $(f(x))^2$ over B_0 can be converted to a matrix multiplication, as shown below:

$$\int (f(x))^2 dx = \int \left[\sum_{z=1}^m \frac{w_z I(x \in G_z)}{|G_z|} \right]^2 dx \\ = \int \sum_{i=1}^m \sum_{j=1}^m \frac{w_i w_j}{|G_i| |G_j|} I(x \in G_i) I(x \in G_j) dx$$

which can be simplified to

$$\begin{aligned} & \sum_{i=1}^m \sum_{j=1}^m \frac{w_i w_j}{|G_i||G_j|} |G_i \cap G_j| \\ &= \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_m \end{bmatrix}^T \begin{bmatrix} \frac{|G_1 \cap G_1|}{|G_1||G_1|} & \cdots & \frac{|G_1 \cap G_m|}{|G_1||G_m|} \\ \vdots & \ddots & \vdots \\ \frac{|G_m \cap G_1|}{|G_m||G_1|} & \cdots & \frac{|G_m \cap G_m|}{|G_m||G_m|} \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_m \end{bmatrix} \\ &= \mathbf{w}^T Q \mathbf{w} \end{aligned}$$

Second, we express the equality constraints in an alternative form. Note that the left-hand side of each equality constraint, i.e., $\int_{B_i} f(x) dx$, can be expressed as:

$$\begin{aligned} \int_{B_i} f(x) dx &= \int_{B_i} \sum_{j=1}^m \frac{w_j}{|G_j|} I(x \in G_j) dx \\ &= \sum_{j=1}^m \frac{w_j}{|G_j|} \int_{B_i} I(x \in G_j) dx = \sum_{j=1}^m \frac{w_j}{|G_j|} |B_i \cap G_j| \\ &= \begin{bmatrix} \frac{|B_i \cap G_1|}{|G_1|} & \cdots & \frac{|B_i \cap G_m|}{|G_m|} \end{bmatrix} \begin{bmatrix} w_1 \\ \vdots \\ w_m \end{bmatrix} \end{aligned}$$

Then, the equality constraints, i.e., $\int_{B_i} f(x) dx = s_i$ for $i = 1, \dots, n$, can be expressed as follows:

$$\begin{aligned} & \begin{bmatrix} \frac{|B_1 \cap G_1|}{|G_1|} & \cdots & \frac{|B_1 \cap G_m|}{|G_m|} \\ \vdots & \ddots & \vdots \\ \frac{|B_n \cap G_1|}{|G_1|} & \cdots & \frac{|B_n \cap G_m|}{|G_m|} \end{bmatrix} \begin{bmatrix} w_1 \\ \vdots \\ w_m \end{bmatrix} = \begin{bmatrix} s_1 \\ \vdots \\ s_m \end{bmatrix} \\ & \Rightarrow A \mathbf{w} = \mathbf{s} \end{aligned}$$

Finally, $\mathbf{w}^T \mathbf{1} = 1$ if and only if $\int f(x) = 1$, and $\mathbf{w} \succeq \mathbf{0}$ for arbitrary G_z if and only if $\int f(x) \geq 0$. \square

The implication of the above theorem is significant: we could reduce the problem of $O(2^n)$ complexity to the problem of only $O(n^2)$ complexity.

Conversion Two: Moving Constraints The quadratic programming problem in Theorem 1 can be solved efficiently by most off-the-shelf optimization libraries; however, we can solve the problem even faster by converting the problem to an alternative form. We first present the alternative problem, then discuss it.

Problem 3 (QuickSel’s QP) *QuickSel solves this problem alternative to the quadratic programming problem in Theorem 1:*

$$\arg \min_{\mathbf{w}} \ell(\mathbf{w}) = \mathbf{w}^T Q \mathbf{w} + \lambda \|\mathbf{A} \mathbf{w} - \mathbf{s}\|^2$$

where λ is a large real value (QuickSel uses $\lambda = 10^6$).

In formulating Problem 3, two types of conversions are performed: (1) the consistency with the observed queries (i.e., $\mathbf{A} \mathbf{w} = \mathbf{s}$) is moved into the optimization objective as a penalty term, and (2) the positivity of $f(x)$ is not explicitly specified (by $\mathbf{w} \succeq \mathbf{0}$). These two types of conversions have little impact on the solution for two reasons. First, to guarantee the consistency, a large penalty (i.e., $\lambda = 10^6$) is used. Second, the mixture model $f(x)$ is bound to approximate the true distribution, which is always non-negative. We empirically examine the advantage of solving Problem 3 (instead of solving the problem in Theorem 1 directly) in §5.7.

The solution \mathbf{w}^* to Problem 3 can be obtained in a straightforward way by setting its gradients of the objective (with respect to \mathbf{w}) equal to $\mathbf{0}$:

$$\begin{aligned} \frac{\partial \ell(\mathbf{w}^*)}{\partial \mathbf{w}} &= 2Q \mathbf{w}^* + 2\lambda A^T (\mathbf{A} \mathbf{w}^* - \mathbf{s}) = \mathbf{0} \\ \Rightarrow \mathbf{w}^* &= (Q + \lambda A^T A)^{-1} \lambda A \mathbf{s} \end{aligned}$$

Observe that \mathbf{w}^* is expressed in a closed form; thus, we can obtain \mathbf{w}^* analytically instead of using iterative procedures typically required for general quadratic programming.

5 EXPERIMENT

In this section, we empirically study QuickSel. In summary, our results show the following:

- End-to-end comparison against other query-driven methods:** QuickSel was significantly faster (34.0×–179.4×) for the same accuracy—and produced much more accurate estimates (26.8%–91.8% lower error) for the same time limit—than previous query-driven methods. (§5.2)
- Comparison against periodic database scans:** For the same storage size, QuickSel’s selectivity estimates were 77.7% and 91.3% more accurate than scan-based histograms and sampling, respectively. (§5.3)
- Impact on PostgreSQL performance:** Using QuickSel for PostgreSQL makes the system 2.25× faster (median) than the default. (§5.4)
- Effectiveness of QuickSel’s mixture model:** QuickSel’s model produced considerably more accurate estimates than histograms given the same number of parameters. (§5.5)
- Robustness to workload shifts:** QuickSel’s accuracy quickly recovers after sudden workload shifts. (§5.6)
- Optimization efficiency:** QuickSel’s optimization method (Problem 3) was 1.5×–17.2× faster than solving the standard quadratic programming. (§5.7)

5.1 Experimental Setup

Methods Our experiments compare QuickSel to six other selectivity estimation methods.

Query-driven Methods:

1. STHoles [19]: This method creates histogram buckets by partitioning existing buckets (as in Figure 1). The frequency of an existing bucket is distributed uniformly among the newly created buckets.
2. ISOMER [93]: This method applies STHoles for histogram bucket creations, but it computes the optimal frequencies of the buckets by finding the maximum entropy distribution. Among existing query-driven methods, ISOMER produced the highest accuracy in our experiments.
3. ISOMER+QP: This method combines ISOMER’s approach for creating histogram buckets and QuickSel’s quadratic programming (Problem 3) for computing the optimal bucket frequencies.
4. QueryModel [13]: This method computes the selectivity estimate by a weighted average of the selectivities of observed queries. The weights are determined based on the similarity of the new query and each of the queries observed in the past.

Scan-based Methods:

5. AutoHist: This method creates an equiwidth multidimensional histogram by scanning the data. It also updates its histogram whenever more than 20% of the data changes (this is the default setting with SQL Server’s AUTO_UPDATE_STATISTICS option [8]).
6. AutoSample: This method relies on a uniform random sample of data to estimate selectivities. Similar to AutoHist, AutoSample updates its sample whenever more than 10% of the data changes.

We have implemented all methods in Java.

Datasets and Query Sets We use two real datasets and one synthetic dataset in our experiments, as follows:

1. DMV: This dataset contains the vehicle registration records of New York State [95]. It contains 11,944,194 rows. Here, the queries ask for the number of valid registrations for vehicles produced within a certain date range. Answering these queries involves predicates on three attributes: `model_year`, `registration_date`, and `expiration_date`.
2. Instacart: This dataset contains the sales records of an online grocery store [94]. We use their orders table, which contains 3.4 million sales records. Here, the queries ask for the reorder frequency for orders made during different hours of the day. Answering these queries involves predicates on two attributes: `order_hour_of_day` and `days_since_prior`. (In §5.3, we use more attributes (up to ten).)
3. Gaussian: We also generated a synthetic dataset using a bivariate dimensional normal distribution. We varied this dataset to study our method under workload shifts, different degrees of correlation between the attributes, and

more. Here, the queries count the number of points that lie within a randomly generated rectangle.

For each dataset, we measured the estimation quality using 100 test queries not used for training. The ranges for selection predicates (in queries) were generated randomly within a feasible region; the ranges of different queries may or may not overlap.

Environment All our experiments were performed on `m5.4xlarge` EC2 instances, with 16-core Intel Xeon 2.5GHz and 64 GB of memory running Ubuntu 16.04.

Metrics We use the root mean square (RMS) error:

$$\text{RMS error} = \left(\frac{1}{t} \sum_{i=1}^t (\text{true_sel} - \text{est_sel})^2 \right)^{1/2}$$

where t is the number of test queries. We report the RMS errors in percentage (by treating both `true_sel` and `est_sel` as percentages).

When reporting training time, we include the time required for refining a model using an additional observed query, which itself includes the time to store the query and run the necessary optimization routines.

5.2 Selectivity Estimation Quality

In this section, we compare the end-to-end selectivity estimation quality of QuickSel versus query-driven histograms. Specifically, we gradually increased the number of observed queries provided to each method from 10 to 1,000. For each number of observed queries, we measured the estimation error and training time of each method using 100 test queries.

These results are reported in Figure 3. Given the same number of observed queries, QuickSel’s training was significantly faster (Figures 3a and 3d), while still achieving comparable estimation errors (Figures 3b and 3e). We also studied the relationship between errors and training times in Figures 3c and 3f, confirming QuickSel’s superior efficiency (STHoles, ISOMER+QP, and QueryModel are omitted in these figures due to their poor performance). In summary, QuickSel was able to quickly learn from a large number of observed queries (i.e., shorter training time) and produce highly accurate models.

5.3 Comparison to Scan-based Methods

We also compared QuickSel to two automatically-updating scan-based methods, AutoHist and AutoSample, which incorporate SQL Server’s automatic updating rule into equiwidth multidimensional histograms and samples, respectively. Since both methods incur an up-front cost for obtaining their statistics, they should produce relatively more accurate estimates initially (before seeing new queries). In

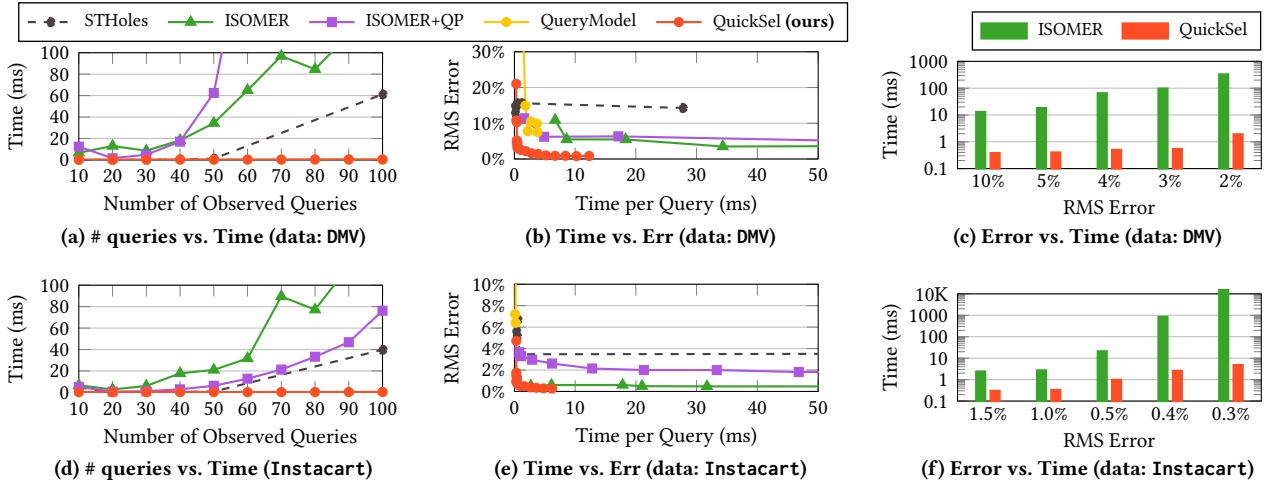


Figure 3: Comparison between QuickSel and query-driven histograms. The lower, the better. Left: The per-query overhead of QuickSel was extremely low. Middle: QuickSel was the most accurate for the same time budget. Right: QuickSel required significantly less time for the same accuracy.




Figure 4: QuickSel versus periodically updating scan-based methods (given the same storage size).

contrast, the accuracy of QuickSel’s estimates should quickly improve as new queries are observed.


To verify this empirically, we first generated a Gaussian dataset (1 million tuples) with correlation 0. We then inserted 200K new tuples generated from a distribution with a *different* correlation after processing 200 queries, and repeated this process. In other words, after processing the first 100 queries, we inserted new data with correlation 0.1; after processing the next 100 queries, we inserted new data with correlation 0.2; and continued this process until a total of 1000 queries were processed. We performed this process for each method under comparison. QuickSel adjusted its model each time after observing 100 queries. AutoHist and AutoSample updated their statistics after each batch of data insertion. QuickSel and AutoHist both used 100 parameters (# of subpopulations for the mixture model and # of buckets for histograms); AutoSample used a sample of 100 tuples.

Figure 4a shows the error of each method. As expected, AutoHist produced more accurate estimates initially. However, as more queries were processed, the error of QuickSel drastically decreased. In contrast, the errors of AutoSample

and AutoHist did not improve with more queries, as they only depend on the frequency at which a new scan (or sampling) is performed. After processing only 100 queries (i.e., initial update), QuickSel produced more accurate estimates than both AutoHist and AutoSample. On average (including the first 100 queries), QuickSel was 71.4% and 89.8% more accurate than AutoHist and AutoSample, respectively. This is consistent with the previously reported observations that query-driven proposals yield better accuracy than scan-based ones [19]. (The reason why query-driven proposals have not been widely adopted to date is due to their prohibitive cost; see §7.2).

In addition, Figure 4b compares the update times of the three methods. By avoiding scans, QuickSel’s query-driven updates were 525× and 243× faster than AutoHist and AutoSample, respectively.

Finally, we studied how the performance of those methods changed as we increased the data dimension (i.e., the number of attributes appearing in selection predicates). First, using the Instacart dataset, we designed each query to target a




(a) Individual Query Speedup over Postgres Default


Speedup Stat	Value
Max	3.47×
Median	2.25×
Average	2.09×
Min	0.98×

(b) Summary of Speedup

Figure 5: QuickSel’s impact on PostgreSQL query performance. We compared (1) PostgreSQL default and (2) PostgreSQL with QuickSel. The left figure shows individual query speedups (note: the original latencies were between 5.2–9.1 secs). The speedup values around 1× were due to no query plan change despite different estimates. The right figure summarizes those speedup numbers.



(a) # of queries vs. # of parameters (data: Instacart)



(b) # of parameters vs. Error (data: Instacart)

Figure 6: Comparison between QuickSel’s model and the models of query-driven histograms. The lower, the better, Left: For the same number of observed queries, QuickSel used the least number of model parameters. Right: QuickSel’s model was more effective in expressing the data distribution, yielding the lowest error.

random subset of dimensions ($N/2$) as increasing the dimension N from 2 to 10. In all test cases (Figure 4c left), QuickSel’s accuracy was consistent, showing its ability to scale to high-dimensional data. Also in this experiment, QuickSel performed significantly better than, or comparably to, histograms and sampling. We could also obtain a similar result using the Gaussian dataset (Figure 4c right). This consistent performance across different data dimensions is primarily due to how QuickSel is designed; that is, its estimation only depends on how much queries overlap with one another.

5.4 Impact on Query Performance

This section examines QuickSel’s impact on query performance. That is, we test if QuickSel’s more accurate selectivity estimates can lead to improved query performance for actual database systems (i.e., shorter latency).

To measure the actual query latencies, we used PostgreSQL ver. 10 with a third-party extension, called `pg_hint_plan` [6]. Using this extension, we enforced our own estimates (for PostgreSQL’s query optimization) in place of the default ones. We compared PostgreSQL Default (i.e., no hint) and QuickSel—to measure the latencies of the following join query in processing the Instacart dataset:

```
select count(*)
from S inner join T on S.tid = T.tid
```

```
inner join U on T.uid = U.uid
where (range_filter_on_T)
and (range_filter_on_S);
```

where the joins keys for the tables S , T , and U were in the PK-FK relationship, as described by the schema (of Instacart).

Figure 5 shows the speedups QuickSel could achieve in comparison to PostgreSQL Default. Note that QuickSel does not improve any underlying I/O or computation speed; its speedups are purely from helping PostgreSQL’s query optimizer choose a more optimal plan based on improved selectivity estimates. Even so, QuickSel could bring 2.25× median speedup, with 3.47× max speedup. In the worst case, PostgreSQL with QuickSel was almost identical to PostgreSQL Default (i.e., 0.98× speedup).

5.5 QuickSel’s Model Effectiveness

In this section, we compare the effectiveness of QuickSel’s model to that of models used in the previous work. Specifically, the effectiveness is assessed by (1) how the model size—its number of parameters—grows as the number of observed queries grows, and (2) how quickly its error decreases as its number of parameters grows.

Figure 6a reports the relationship between the number of observed queries and the number of model parameters. As discussed in §2.3, the number of buckets (hence, parameters)




Figure 7: Robustness to sudden workload shifts, which occurred at the sequence #100 and at #200. QuickSel’s error increased temporarily right after each workload jump, but it reduced soon.




Figure 8: QuickSel’s optimization effect.

of ISOMER increased quickly as the number of observed queries grew. STHoles was able to keep the number of its parameters small due to its bucket merging technique; however, this had a negative impact on its accuracy. Here, QuickSel used the least number of model parameters. For instance, when 100 queries were observed for DMV, QuickSel had 10× fewer parameters than STHoles and 56× fewer parameters than ISOMER.

We also studied the relationship between the number of model parameters and the error. The lower the error (for the same number of model parameters), the more effective the model. Figure 6b shows the result. Given the same number of model parameters, QuickSel produced significantly more accurate estimates. Equivalently, QuickSel produced the same quality estimates with much fewer model parameters.

5.6 Robustness to Workload Shifts

In this section, we test QuickSel’s performance under significant workload shifts. That is, after observing a certain number of queries (i.e., 100 queries) around a certain region of data, the query workload suddenly jumps to a novel region. This pattern repeats several times.

Figure 7 shows the result. Here, we could observe the following pattern. QuickSel’s error increased significantly right after each jump (i.e., at query sequence #100 and at #200), producing 1.5×–3.6× higher RMS errors compared to histograms. However, QuickSel’s error dropped quickly, achieving 12×–378× lower RMS errors than histograms. This was possible due to QuickSel’s faster adaptation.

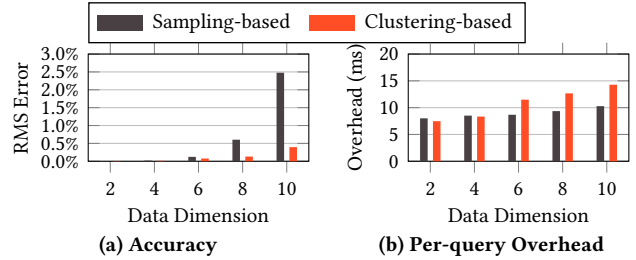


Figure 9: Subpopulation generation approaches. Clustering-based was more accurate, but slower.

5.7 QuickSel Internal Mechanisms

In this section, we empirically study (1) the effect of QuickSel’s optimization (presented in §4.2), and (2) two alternative mechanisms for generating subpopulations (presented in §3.3).

Optimization Efficiency To study QuickSel’s optimization efficiency, we compared two approaches for solving the quadratic problem defined in Theorem 1: solving the original QP without any modifications versus solving our modified version (Problem 3). We used the `cvxopt` library for the former and used `joblas` (a linear algebra library) for the latter. Both libraries use multiple cores for parallel processing.

Figure 8 shows the time taken by each optimization approach. The second approach (Problem 3) was increasingly more efficient as the number of observed queries grew. For example, it was 8.36× faster when the number of observed queries reached 1,000. This is thanks to the modified problem having an analytical solution, while the original problem required an iterative gradient descent solution.

Subpopulation Generation We empirically studied the two subpopulation generation approaches (i.e., the sampling-based approach and the clustering-based approach, §3.3) in terms of their scalability to high-dimensional data. Specifically, we compared their estimation accuracies and computational overhead using the Gaussian dataset (with its dimension set to 2–10).

Figure 9 reports the results. As shown in Figure 9(a), the clustering-based approach obtained impressive accuracy in comparison to the sampling-based one. However, as shown in Figure 9(b), the clustering-based approach produced higher overhead (i.e., longer training times), which is an example of the natural tradeoff between cost and accuracy.

6 CONNECTION: MSE AND ENTROPY

The max-entropy query-driven histograms optimize their parameters (i.e., bucket frequencies) by searching for the parameter values that maximize the entropy of the distribution $f(x)$. We show that this approach is approximately

Table 3: Comparison of selectivity estimation methods

Approach	Model	Method	Key Contributions
Based on Database Scans (Scan-based Selectivity Estimation)	Histograms	Multi-dim Hist [27, 42, 70]	Introduces multidimensional histograms
		Muralikrishna [79]	Introduces equidepth histograms
		Van Gelder [106]	Estimates Join selectivity with histograms for <i>important</i> domains
		GOH [44]	Optimizes single-attribute histograms for joint distribution
		Thaper [101]	Builds histograms over streaming data
	Sampling	To [102]	Builds histograms with entropy as a metric
		Lipton [68]	Introduces adaptive sampling for high accuracy
		Haas [37]	Uses sampling for join selectivity estimation
	ML	Riondato [87]	Guarantees accuracy relying on the VC-dimension of queries
KDE [34, 35, 40]		Applies kernel density estimation to selectivity estimation	
PGM [32, 92, 104]		Uses probabilistic graphical models for selectivity estimation	
Based on Observed Queries (Query-driven Selectivity Estimation)	Error-feedback Histograms (<i>fast but less accurate</i>)	Neural Net [56, 69]	Trains a neural network for selectivity estimation
		ST-histogram [11]	Refines the bucket frequencies based on the errors
		LEO [96]	Identifies incorrect statistics using observed queries
		STHoles [19]	Proposes a new buckets split mechanism; adopted by ISOMER
	Max-Entropy Histograms (<i>accurate but slow</i>)	SASH [66]	Proposes a <i>junction tree</i> model for finding the best set of histograms
		QueryModel [13]	Avoids modeling the data distribution by using queries directly
		ISOMER [74, 75, 93]	Finds a maximum entropy distribution consistent with observed queries
	Mixture Model (<i>fast & accurate</i>)	Kaushik et al. [52]	Extends ISOMER for distinct values
		Ré et al. [85, 86]	Seeks the max entropy distribution based on <i>possible worlds</i>
	QuickSel (Ours)	Employs a mixture model for selectivity estimation; develops an efficient training algorithm for the new model	

equivalent to QuickSel’s optimization objective, i.e., minimizing the mean squared error (MSE) of $f(x)$ from a uniform distribution. The entropy of the probability density function is defined as $-\int f(x) \log(f(x)) dx$. Thus, maximizing the entropy is equivalent to minimizing $\int f(x) \log(f(x)) dx$, which is related to minimizing MSE as follows:

$$\begin{aligned} \arg \min \int f(x) \log(f(x)) dx &\approx \arg \min \int f(x) (f(x) - 1) dx \\ &= \arg \min \int (f(x))^2 dx \end{aligned}$$

since $\int f(x) dx = 1$ by definition. We used the first-order Taylor expansion to approximate $\log(x)$ with $x - 1$. Note that, when the constraint $\int f(x) dx = 1$ is considered, $f(x) = 1/|R_0|$ is the common solution to both the entropy maximization and minimizing MSE.

7 RELATED WORK

There is extensive work on selectivity estimation due to its importance for query optimization. In this section, we review both scan-based (§7.1) and query-driven methods (§7.2). QuickSel belongs to the latter category. We have summarized the related work in Table 3.

7.1 Database Scan-based Estimation

As explained in §1, we use the term *scan-based methods* to refer to techniques that directly inspect the data (or part of it) for collecting their statistics. These approaches differ from query-based methods which rely only on the actual selectivities of the observed queries.

Scan-based Histograms These approaches approximate the joint distribution by periodically scanning the data. There has been much work on how to efficiently express the joint distribution of multidimensional data [23, 25, 27, 33–35, 37, 39, 40, 42–44, 49, 51, 64, 68, 70, 77, 79, 87, 101, 102, 106]. There is also some work on histograms for special types of data, such as XML [10, 16, 107, 109], spatial data [48, 57–59, 67, 72, 80, 97, 99, 100, 108, 112, 113], graph [29], string [45–47, 76]; or for privacy [38, 63, 65].

Sampling Sampling-based methods rely on a sample of data for estimating its joint distribution [37, 68, 87]. However, drawing a new random sample requires a table-scan or random retrieval of tuples, both of which are costly operations and hence, are only performed periodically.

Machine Learning Models KDE is a technique that translates randomly sampled data points into a distribution [91]. In the context of selectivity estimation, KDE has been used as an alternative to histograms [34, 35, 40]. Like mixture models (MM), KDE also expresses a probability density function as a summation of some basis functions. However, KDE and MM are fundamentally different. KDE relies on independent and identically distributed samples, and hence lends itself to scan-based selectivity estimation. In contrast, MM does not require any sampling and can thus be used in query-driven selectivity estimation (where sampling is not practical). Similarly, probabilistic graphical models [32, 92, 104], neural networks [56, 69], and tree-based ensembles [28] have been used for selectivity estimation. Unlike histograms, these approaches can capture column correlations more succinctly. However, applicability of these models for query-driven selectivity estimation has not been explored and remains unclear.

More recently, sketching [20] and probe executions [103] have been proposed, which differ from ours in that they build their models directly using the data (not query results). Similar to histograms, using the data requires either periodic updates or higher query processing overhead. QuickSel avoids both of these shortcomings with its query-driven MM.

7.2 Query-driven Estimation

Query-driven techniques create their histogram buckets adaptively according to the queries they observe in the workload. These can be further categorized into two techniques based on how they compute their bucket frequencies: error-feedback histograms and max-entropy histograms.

Error-feedback Histograms Error-feedback histograms [11, 13, 19, 54, 55, 66] adjust bucket frequencies in consideration of the errors made by old bucket frequencies. They differ in how they create histogram buckets according to the observed queries. For example, STHoles [19] splits existing buckets with the predicate range of the new query. SASH [66] uses a space-efficient multidimensional histogram, called MHIST [27], but determines its bucket frequencies with an error-feedback mechanism. QueryModel [13] treats the observed queries themselves as conceptual histogram buckets and determines the distances among those buckets based on the similarities among the queries’ predicates.

Max-Entropy Histograms Max-entropy histograms [52, 74, 75, 86, 93] find a maximum entropy distribution consistent with the observed queries. Unfortunately, these methods generally suffer from the exponential growth in their number of buckets as the number of observed queries grows (as discussed in §2). QuickSel avoids this problem by relying on mixture models.

Fitting Parametric Functions Adaptive selectivity estimation [22] fits a parametric function (e.g., linear, polynomial) to the observed queries. This approach is more applicable when we know the data distribution *a priori*, which is not assumed by QuickSel.

Self-tuning Databases Query-driven histograms have also been studied in the context of self-tuning databases [60, 71, 73]. IBM’s LEO [96] corrects errors in any stage of query execution based on the observed queries. Microsoft’s AutoAdmin [12, 21] focuses on automatic physical design, self-tuning histograms, and monitoring infrastructure. Part of this effort is ST-histogram [11] and STHoles [19] (see Table 3). DBL [81] and IDEA [30] exploit the answers to past queries for more accurate approximate query processing. QueryBot 5000 [71] forecasts the future queries, whereas OtterTune [105] and index [61] use machine learning for automatic physical design and building secondary indices, respectively.

8 CONCLUSION AND FUTURE WORK

The prohibitive cost of query-driven selectivity estimation techniques has greatly limited their adoption by DBMS vendors, which for the most part still rely on scan-based histograms and samples that are periodically updated and are otherwise stale. In this paper, we proposed a new framework, called *selectivity learning* or QuickSel, which learns from every query to continuously refine its internal model of the underlying data, and therefore produce increasingly more accurate selectivity estimates over time. QuickSel differs from previous query-driven selectivity estimation techniques by (i) not using histograms and (ii) enabling extremely fast refinements using its mixture model. We formally showed that the training cost of our mixture model can be reduced from exponential to only quadratic complexity (Theorem 1).

Supporting Complex Joins When modeling the selectivity of join queries, even state-of-the-art modes [28, 56, 111] take a relatively simple approach: conceptually prejoining corresponding tables and constructing a joint probability distribution over each join pattern. We plan to similarly extend our current formulation of QuickSel to model the selectivity of general joins.

9 ACKNOWLEDGEMENT

This material is based upon work supported by the National Science Foundation under Grant No. 1629397 and the Michigan Institute for Data Science (MIDAS) PODS. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] Apache commons: Optimization. <http://commons.apache.org/proper/commons-math/userguide/optimization.html>. [Online; accessed September-16-2018].
- [2] Collecting histogram statistics. https://www.ibm.com/support/knowledgecenter/SSEPEK_11.0.0/perf/src/tpc/db2z_collecthistogramstatistics.html. [Online; accessed September-16-2018].
- [3] Histogram-based statistics. <https://mariadb.com/kb/en/library/histogram-based-statistics/>. [Online; accessed September-16-2018].
- [4] Joptimizer. <http://www.joptimizer.com/>. [Online; accessed September-16-2018].
- [5] MATLAB: quadprog. <https://www.mathworks.com/help/optim/ug/quadprog.html>. [Online; accessed September-16-2018].
- [6] pg_hint_plan 1.1. https://pghintplan.osdn.jp/pg_hint_plan.html. [Online; accessed February-13-2020].
- [7] PostgreSQL 9.2.24 documentation. <https://www.postgresql.org/docs/9.2/static/row-estimation-examples.html>. [Online; accessed September-16-2018].
- [8] Statistical maintenance functionality (autostats) in sql server. <https://support.microsoft.com/en-us/help/195565/statistical-maintenance-functionality-autostats-in-sql-server>. [Online; accessed September-16-2018].
- [9] Statistics. <https://docs.microsoft.com/en-us/sql/relational-databases/statistics/statistics?view=sql-server-2017>. [Online; accessed September-16-2018].
- [10] A. Abounaga, A. R. Alameldeen, and J. F. Naughton. Estimating the selectivity of xml path expressions for internet scale applications. In *VLDB*, 2001.
- [11] A. Abounaga and S. Chaudhuri. Self-tuning histograms: Building histograms without looking at data. *SIGMOD*, 1999.
- [12] S. Agrawal, N. Bruno, S. Chaudhuri, and V. R. Narasayya. Autoadmin: Self-tuning database system technology. *IEEE Data Eng. Bull.*, 2006.
- [13] C. Anagnostopoulos and P. Triantafyllou. Learning to accurately count with query-driven predictive analytics. In *Big Data*, 2015.
- [14] M. Andersen, J. Dahl, and L. Vandenberghe. Cvxopt: A python package for convex optimization. 2013.
- [15] T. Asparouhov and B. Muthén. Structural equation models and mixture models with continuous nonnormal skewed distributions. *Structural Equation Modeling: A Multidisciplinary Journal*, 2016.
- [16] S. S. Bhowmick, E. Leonardi, and H. Sun. Efficient evaluation of high-selective xml twig patterns with parent child edges in tree-unaware rdbs. In *SIGMOD*, 2007.
- [17] C. M. Bishop. *Pattern recognition and machine learning*. 2006.
- [18] B. Blohsfeld, D. Korus, and B. Seeger. A comparison of selectivity estimators for range queries on metric attributes. In *SIGMOD Record*, 1999.
- [19] N. Bruno, S. Chaudhuri, and L. Gravano. Stholes: a multidimensional workload-aware histogram. In *SIGMOD*, 2001.
- [20] W. Cai, M. Balazinska, and D. Suciu. Pessimistic cardinality estimation: Tighter upper bounds for intermediate join cardinalities. In *SIGMOD*, 2019.
- [21] S. Chaudhuri and V. Narasayya. Self-tuning database systems: a decade of progress. In *PVLDB*, 2007.
- [22] C. M. Chen and N. Roussopoulos. Adaptive selectivity estimation using query feedback. In *SIGMOD*, 1994.
- [23] L. Chen and M. T. Ozsu. Multi-scale histograms for answering queries over time series data. In *ICDE*, 2004.
- [24] Y.-C. Chen. Lecture 6: Density estimation: Histogram and kernel density estimator. http://faculty.washington.edu/yenchic/18W_425/Lec6_hist_KDE.pdf. [Online; accessed September-16-2018].
- [25] G. Cormode, M. Garofalakis, P. J. Haas, C. Jermaine, et al. Synopses for massive data: Samples, histograms, wavelets, sketches. 2011.
- [26] P. F. Craigmile and D. Tirrerington. Parameter estimation for finite mixtures of uniform distributions. *Communications in Statistics-Theory and Methods*, 1997.
- [27] A. Deshpande, M. Garofalakis, and R. Rastogi. Independence is good: Dependency-based histogram synopses for high-dimensional data. *SIGMOD Record*, 2001.
- [28] A. Dutt, C. Wang, A. Nazi, S. Kandula, V. Narasayya, and S. Chaudhuri. Selectivity estimation for range predicates using lightweight models. *PVLDB*, 2019.
- [29] J. Feng, Q. Qian, Y. Liao, G. Li, and N. Ta. Dmt: a flexible and versatile selectivity estimation approach for graph query. In *WAIM*, 2005.
- [30] A. Galakatos, A. Crotty, E. Zraggen, C. Binnig, and T. Kraska. Revisiting reuse for approximate query processing. *PVLDB*, 2017.
- [31] A. Genz. Numerical computation of multivariate normal probabilities. *Journal of computational and graphical statistics*, 1992.
- [32] L. Getoor, B. Taskar, and D. Koller. Selectivity estimation using probabilistic models. In *ACM SIGMOD Record*, volume 30, pages 461–472. ACM, 2001.
- [33] S. Guha, N. Koudas, and D. Srivastava. Fast algorithms for hierarchical range histogram construction. In *PODS*, 2002.
- [34] D. Gunopulos, G. Kollios, V. J. Tsotras, and C. Domeniconi. Approximating multi-dimensional aggregate range queries over real attributes. In *SIGMOD Record*, 2000.
- [35] D. Gunopulos, G. Kollios, V. J. Tsotras, and C. Domeniconi. Selectivity estimators for multidimensional range queries over real attributes. *VLDBJ*, 2005.
- [36] A. Gupta and T. Miyawaki. On a uniform mixture model. *Biometrical Journal*, 1978.
- [37] P. J. Haas, J. F. Naughton, and A. N. Swami. On the relative cost of sampling for join selectivity estimation. In *PODS*, 1994.
- [38] M. Hay, V. Rastogi, G. Miklau, and D. Suciu. Boosting the accuracy of differentially private histograms through consistency. *PVLDB*, 2010.
- [39] Z. He, X. Xu, S. Deng, and B. Dong. K-histograms: An efficient clustering algorithm for categorical dataset. *arXiv preprint cs/0509033*, 2005.
- [40] M. Heimeel, M. Kiefer, and V. Markl. Self-tuning, gpu-accelerated kernel density models for multidimensional selectivity estimation. In *SIGMOD*, 2015.
- [41] R. Hu, Z. Wang, W. Fan, and S. Agarwal. Cost based optimizer in apache spark 2.2. <https://databricks.com/blog/2017/08/31/cost-based-optimizer-in-apache-spark-2-2.html>, 2018. [Online; accessed September-16-2018].
- [42] I. F. Ilyas, V. Markl, P. Haas, P. Brown, and A. Abounaga. Cords: automatic discovery of correlations and soft functional dependencies. In *SIGMOD*, 2004.
- [43] Z. Istvan, L. Woods, and G. Alonso. Histograms as a side effect of data movement for big data. In *SIGMOD*, 2014.
- [44] H. Jagadish, H. Jin, B. C. Ooi, and K.-L. Tan. Global optimization of histograms. *SIGMOD Record*, 2001.
- [45] H. Jagadish, O. Kapitskaia, R. T. Ng, and D. Srivastava. Multi-dimensional substring selectivity estimation. In *VLDB*, 1999.
- [46] H. Jagadish, O. Kapitskaia, R. T. Ng, and D. Srivastava. One-dimensional and multi-dimensional substring selectivity estimation. *VLDB*, 2000.
- [47] H. Jagadish, R. T. Ng, and D. Srivastava. Substring selectivity estimation. In *PODS*, 1999.
- [48] H. V. Jagadish, N. Koudas, S. Muthukrishnan, V. Poosala, K. C. Sevcik, and T. Suel. Optimal histograms with quality guarantees. In *VLDB*, 1998.
- [49] J. Jests, K. Yi, and F. Li. Building wavelet histograms on large data in mapreduce. *PVLDB*, 2011.

- [50] H. Joe. Approximations to multivariate normal rectangle probabilities based on conditional expectations. *Journal of the American Statistical Association*, 1995.
- [51] P. Karras and N. Mamoulis. Lattice histograms: a resilient synopsis structure. In *ICDE*, 2008.
- [52] R. Kaushik and D. Suci. Consistent histograms in the presence of distinct value counts. *PVLDB*, 2009.
- [53] M. S. Kester, M. Athanassoulis, and S. Idreos. Access path selection in main-memory optimized data systems: Should i scan or should i probe? In *SIGMOD*, 2017.
- [54] A. Khachatryan, E. Müller, C. Stier, and K. Böhm. Sensitivity of self-tuning histograms: query order affecting accuracy and robustness. In *SSDBM*, 2012.
- [55] A. Khachatryan, E. Müller, C. Stier, and K. Böhm. Improving accuracy and robustness of self-tuning histograms by subspace clustering. *TKDE*, 2015.
- [56] A. Kipf, T. Kipf, B. Radke, V. Leis, P. Boncz, and A. Kemper. Learned cardinalities: Estimating correlated joins with deep learning. *arXiv preprint arXiv:1809.00677*, 2018.
- [57] G. Koloniari, Y. Petrakis, E. Pitoura, and T. Tsotsos. Query workload-aware overlay construction using histograms. In *CIKM*, 2005.
- [58] F. Korn, T. Johnson, and H. Jagadish. Range selectivity estimation for continuous attributes. In *ssdbm*, 1999.
- [59] N. Koudas, S. Muthukrishnan, and D. Srivastava. Optimal histograms for hierarchical range queries. In *PODS*, 2000.
- [60] T. Kraska, M. Alizadeh, A. Beutel, E. H. Chi, A. Kristo, G. Leclerc, S. Madden, H. Mao, and V. Nathan. Sagedb: A learned database system. In *CIDR*, 2019.
- [61] T. Kraska, A. Beutel, E. H. Chi, J. Dean, and N. Polyzotis. The case for learned index structures. In *SIGMOD*, 2018.
- [62] S. Kullback and R. A. Leibler. On information and sufficiency. *The annals of mathematical statistics*, 1951.
- [63] Y.-H. Kuo, C.-C. Chiu, D. Kifer, M. Hay, and A. Machanavajjhala. Differentially private hierarchical count-of-counts histograms. *PVLDB*, 2018.
- [64] E. Lam and K. Salem. Dynamic histograms for non-stationary updates. In *IDEAS*, 2005.
- [65] C. Li, M. Hay, V. Rastogi, G. Miklau, and A. McGregor. Optimizing linear counting queries under differential privacy. In *PODS*, 2010.
- [66] L. Lim, M. Wang, and J. S. Vitter. Sash: A self-adaptive histogram set for dynamically changing workloads. In *VLDB*, 2003.
- [67] X. Lin, Q. Liu, Y. Yuan, and X. Zhou. Multiscale histograms: Summarizing topological relations in large spatial datasets. In *VLDB*, 2003.
- [68] R. J. Lipton, J. F. Naughton, and D. A. Schneider. *Practical selectivity estimation through adaptive sampling*. 1990.
- [69] H. Liu, M. Xu, Z. Yu, V. Corvinelli, and C. Zuzarte. Cardinality estimation using neural networks. In *Proceedings of the 25th Annual International Conference on Computer Science and Software Engineering*, CASCON '15, pages 53–59, Riverton, NJ, USA, 2015. IBM Corp.
- [70] C. A. Lynch. Selectivity estimation and query optimization in large databases with highly skewed distribution of column values. In *VLDB*, 1988.
- [71] L. Ma, D. Van Aken, A. Hefny, G. Mezerhane, A. Pavlo, and G. J. Gordon. Query-based workload forecasting for self-driving database management systems. In *SIGMOD*, 2018.
- [72] N. Mamoulis and D. Papadias. Selectivity estimation of complex spatial queries. In *International Symposium on Spatial and Temporal Databases*, pages 155–174, 2001.
- [73] R. C. Marcus, P. Negi, H. Mao, C. Zhang, M. Alizadeh, T. Kraska, O. Papaemmanouil, and N. Tatbul. Neo: A learned query optimizer. *PVLDB*, 2019.
- [74] V. Markl, P. J. Haas, M. Kutsch, N. Megiddo, U. Srivastava, and T. M. Tran. Consistent selectivity estimation via maximum entropy. *VLDBJ*, 2007.
- [75] V. Markl, N. Megiddo, M. Kutsch, T. M. Tran, P. Haas, and U. Srivastava. Consistently estimating the selectivity of conjuncts of predicates. In *PVLDB*, 2005.
- [76] A. Mazeika, M. H. Böhlen, N. Koudas, and D. Srivastava. Estimating the selectivity of approximate string queries. *TODS*, 2007.
- [77] G. Moerkotte, D. DeHaan, N. May, A. Nica, and A. Boehm. Exploiting ordered dictionaries to efficiently construct histograms with q-error guarantees in sap hana. In *SIGMOD*, 2014.
- [78] G. Moser, S. H. Lee, B. J. Hayes, M. E. Goddard, N. R. Wray, and P. M. Visscher. Simultaneous discovery, estimation and prediction analysis of complex traits using a bayesian mixture model. *PLoS genetics*, 2015.
- [79] M. Muralikrishna and D. J. DeWitt. Equi-depth multidimensional histograms. In *SIGMOD Record*, 1988.
- [80] T. Neumann and S. Michel. Smooth interpolating histograms with error guarantees. In *British National Conference on Databases*, 2008.
- [81] Y. Park, A. S. Tajik, M. Cafarella, and B. Mozafari. Database learning: Toward a database that becomes smarter every time. In *SIGMOD*, 2017.
- [82] Y. Park, S. Zhong, and B. Mozafari. Extended report of QuickSel: Quick selectivity learning with mixture models. <https://arxiv.org/abs/1812.10568>, 2020.
- [83] C. Proteau. Guide to performance and tuning: Query performance and sampled selectivity. <http://www.oracle.com/technetwork/products/rdb/0403-sampled-selectivity-128646.pdf>. [Online; accessed September-16-2018].
- [84] S. Ragothaman, S. Narasimhan, M. G. Basavaraj, and R. Dewar. Unsupervised segmentation of cervical cell images using gaussian mixture model. In *CVPR Workshops*, 2016.
- [85] C. Ré and D. Suci. Understanding cardinality estimation using entropy maximization. In *PODS*, 2010.
- [86] C. Ré and D. Suci. Understanding cardinality estimation using entropy maximization. *TODS*, 2012.
- [87] M. Riondato, M. Akdere, U. Çetintemel, S. B. Zdonik, and E. Upfal. The vc-dimension of sql queries and selectivity estimation through sampling. In *ECML PKDD*, 2011.
- [88] J. C. Rodríguez. An overview on optimization in apache hive: Past, present future. <https://www.slideshare.net/HadoopSummit/an-overview-on-optimization-in-apache-hive-past-present-future>. [Online; accessed September-16-2018].
- [89] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover's distance as a metric for image retrieval. *International journal of computer vision*, 2000.
- [90] P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price. Access path selection in a relational database management system. In *SIGMOD*, 1979.
- [91] B. W. Silverman. *Density estimation for statistics and data analysis*. 2018.
- [92] J. Spiegel and N. Polyzotis. Graph-based synopses for relational selectivity estimation. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*, SIGMOD '06, pages 205–216, New York, NY, USA, 2006. ACM.
- [93] U. Srivastava, P. J. Haas, V. Markl, M. Kutsch, and T. M. Tran. Isomer: Consistent histogram construction using query feedback. In *ICDE*, 2006.
- [94] J. Stanley. 3 million instacart orders, open sourced. <https://www.instacart.com/datasets/grocery-shopping-2017>, 2017. [Online; accessed September-16-2018].
- [95] State of New York. Vehicle, snowmobile, and boat registrations. <https://catalog.data.gov/dataset/>

- vehicle-snowmobile-and-boat-registrations, 2018. [Online; accessed September-16-2018].
- [96] M. Stillger, G. M. Lohman, V. Markl, and M. Kandil. Leo-db2's learning optimizer. In *VLDB*, 2001.
- [97] J. Sun, Y. Tao, D. Papadias, and G. Kollios. Spatio-temporal join selectivity. *Information Systems*, 2006.
- [98] A. Swami and K. B. Schiefer. On the estimation of join result sizes. In *EDBT*, 1994.
- [99] M. Tang and F. Li. Scalable histograms on large probabilistic data. In *KDD*, 2014.
- [100] Y. Tao, J. Sun, and D. Papadias. Selectivity estimation for predictive spatio-temporal queries. In *ICDE*, 2003.
- [101] N. Thaper, S. Guha, P. Indyk, and N. Koudas. Dynamic multidimensional histograms. In *SIGMOD*, 2002.
- [102] H. To, K. Chiang, and C. Shahabi. Entropy-based histograms for selectivity estimation. In *CIKM*, 2013.
- [103] I. Trummer. Exact cardinality query optimization with bounded execution cost. In *SIGMOD*, 2019.
- [104] K. Tzoumas, A. Deshpande, and C. S. Jensen. Efficiently adapting graphical models for selectivity estimation. *The VLDB Journal*, 22(1):3–27, Feb. 2013.
- [105] D. Van Aken, A. Pavlo, G. J. Gordon, and B. Zhang. Automatic database management system tuning through large-scale machine learning. In *SIGMOD*, 2017.
- [106] A. Van Gelder. Multiple join size estimation by virtual domains. In *PODS*, 1993.
- [107] C. Wang, S. Parthasarathy, and R. Jin. A decomposition-based probabilistic framework for estimating the selectivity of xml twig queries. In *EDBT*, 2006.
- [108] X. Wang, Y. Zhang, W. Zhang, X. Lin, and W. Wang. Selectivity estimation on streaming spatio-textual data using local correlations. *PVLDB*, 2014.
- [109] Y. Wu, J. M. Patel, and H. Jagadish. Using histograms to estimate answer sizes for xml queries. *Information Systems*, 2003.
- [110] J. Yang, X. Liao, X. Yuan, P. Llull, D. J. Brady, G. Sapiro, and L. Carin. Compressive sensing by learning a gaussian mixture model from measurements. *IEEE Transactions on Image Processing*, 2015.
- [111] Z. Yang, E. Liang, A. Kamsetty, C. Wu, Y. Duan, X. Chen, P. Abbeel, J. M. Hellerstein, S. Krishnan, and I. Stoica. Selectivity estimation with deep likelihood models. *CoRR*, abs/1905.04278, 2019.
- [112] Q. Zhang and X. Lin. On linear-spline based histograms. In *WAIM*, 2002.
- [113] Q. Zhang and X. Lin. Clustering moving objects for spatio-temporal selectivity estimation. In *Australasian database conference*, 2004.