# Facilitating SQL Query Composition and Analysis

Zainab Zolaktaf University of British Columbia Vancouver, BC, Canada zolaktaf@cs.ubc.ca Mostafa Milani University of British Columbia Vancouver, BC, Canada mkmilani@cs.ubc.ca Rachel Pottinger University of British Columbia Vancouver, BC, Canada rap@cs.ubc.ca

# ABSTRACT

Formulating efficient SQL queries requires several cycles of tuning and execution. We examine methods that can accelerate and improve this interaction by providing insights about SQL queries *prior to execution*. We achieve this by predicting properties such as the query answer size, its run-time, and error class. Unlike existing approaches, our approach does not rely on any statistics from the database instance or query execution plans. Our approach is based on using data-driven machine learning techniques that rely on large query workloads to model SQL queries and their properties. Empirical results show that the neural network models are more accurate in predicting several query properties.

# **KEYWORDS**

Query Property Prediction, Workload, Machine Learning

#### **ACM Reference Format:**

Zainab Zolaktaf, Mostafa Milani, and Rachel Pottinger. 2020. Facilitating SQL Query Composition and Analysis. 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, 16 pages. https://doi.org/10.1145/3318464.3380602

# **1 INTRODUCTION**

Formulating effective SQL queries is one of the main challenges of interacting with large relational databases. Particularly for inexperienced users, writing good SQL queries may require several cycles of tuning and execution. This diminishes the user experience, prevents them from easily accessing information [20], and can also be costly. For example, cloud-based services, such as Google BigQuery, charge their users for running queries [15, 18, 19]. Moreover, inefficient SQL queries can pose a burden on the database's

*SIGMOD'20, June 14–19, 2020, Portland, OR, USA* © 2020 Association for Computing Machinery. ACM ISBN 978-1-4503-6735-6/20/06...\$15.00 https://doi.org/10.1145/3318464.3380602 resources. Our focus in this work is on facilitating user interaction with the database by providing additional information about SQL queries prior to their execution.

We focus on two groups of users: *end users* and *database administrators (DBAs)*. To help end users compose SQL queries, we study three problems: *query answer size prediction, query run-time prediction*, and *query error prediction*. We can save end users time and effort by pointing out when their queries are inefficient, unlikely to work at all, or are likely to take radically different time than they are expecting (and thus are likely to be not the queries that they are trying to write).

We also improve user interaction for DBAs. To characterize how end users and programs use the service, DBAs need to analyze incoming requests and queries and categorize them into *classes of clients* [47]. This in turn allows them to improve service quality for end users. To help DBAs with this analysis, we study the problem of *session type classification*, which is the automatic identification of the class of clients that generated the queries in a session.

While we use estimates of SQL query properties to improve usability, these estimates have typically been used to improve tasks like admission control, access control, scheduling, and costing during query optimization [1, 13, 38, 39]. Most of these studies, however, are based on manually constructed cost models in the query optimizer and require access to the database instance. But the analytical cost models in the query optimizer can be imprecise due to simplifying assumptions, e.g., uniform data distributions [10, 13, 36]. Moreover, access to the database instance can be infeasible in an increasingly large number of settings, e.g., cloud-based data warehouses like Google BigQuery [15], databases on the hidden web, sources located behind wrappers in data integration systems [5], and instances with limited access due to cost or privacy issues [15, 16, 18, 19]. Due to these restrictions, there is growing need for work that does not assume direct access to database instances.

Our approach for modeling the queries and their properties relies on using *SQL query workloads*, which contain logs of past queries submitted to the database. Query workloads are an alternative resource in settings with limited database instance access. They have been used to improve query performance estimates for tasks like query optimization and scheduling [1, 13, 38, 39]. In addition to requiring access to the database instance and schema, these works examine

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 ACM 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.

synthetic or small-scale query workloads, e.g., TPC-H [52]. They extract hand-engineered features from query execution plans, and apply a prediction model. However, synthetic and small-scale query workloads do not represent the full capacity and challenges of potential queries, as shown empirically in [36].

We use SQL query workloads that are *large-scale and real-world* and present an abundance of realistic usage patterns from a variety of different users. These workloads are broadly used and publicly available in scientific and academic research domains. Examples are the Sloan Digital Sky Survey (SDSS) query workload [44, 45, 50] and SQLShare [22]. Companies and organizations, such as Snowflake [21, 23], often maintain their own large-scale private workloads. DBMSs support logging features that make it easy to generate and maintain these workloads.

When large-scale workloads do not exist for a database, knowledge learned from workloads over other databases may be used for query property prediction. Consequently, we define different query facilitation problem settings that vary in their data heterogeneity. This makes the prediction problem more challenging due to different underlying data distributions for the SQL query and the workload. These settings require models that generalize well and can transfer knowledge learned from workloads to predict query properties over a different database. We empirically show that some models generalize better, which allows reusing large-scale workloads for query performance property prediction. We hope in light of research like ours that shows the benefits of large-scale workloads, more companies and organizations develop, maintain, and share their query workloads, which can ultimately improve transparency and customer engagement.

We use data-driven machine learning techniques that require large-scale workloads to build effective models and predict different query properties. Since there are no standard models for our problems, we start by establishing appropriate baseline models. Our query facilitating problems are essentially query labeling tasks. A closely related area is text categorization, where the goal is to predict categories for documents written in natural language [8, 25, 26, 28, 29, 31, 59]. We chose a broad set of applicable models from this domain; on large-scale datasets, the dominant approaches use Long-Term Short-Term (LSTM) and Convolutional Neural Network (CNN) models [3, 8, 35, 54]. LSTMs treat texts as sequential inputs, while CNNs can automatically identify n-grams. SQL queries have significant differences with natural language sentences, e.g., they include mathematical expressions that are important to retain in the query representation. We therefore applied all models at both character and word level.

Our work is closest to [21, 23], which assume no direct access to database instances and apply data-driven models

SELECT ... FROM PhotoObj
WHERE flags & dbo.fPhotoFlags('BLENDED') > 0

# Figure 1: An inefficient query in which the function in the WHERE clause is called once per matching row.

to large-scale query workloads. They address workload management tasks, such as index recommendation and security audits. We introduce additional problems related to facilitating query composition and analysis, e.g., session classification. Moreover, we formalise and study different problem settings and conduct an extensive workload analysis that results in encouraging empirical results.

In summary we contribute the following:

- We motivate (Section 2) and formally define (Section 3) four problems that help end users and DBAs with composition and analysis by providing query insights prior to execution.
- Our approach relies on exploiting large query workloads. We use two real-world query workloads that are publicly available, SDSS and SQLShare. In Section 4, we describe these workloads. Moreover, to enable better model selection and evaluation, we conduct a comprehensive workload analysis that covers structural and syntactic features of the queries and their labels.
- We examine a broad set of models to establish the baselines and assess the feasibility of our problems in Section 5. We adapted two classes of neural network models to our problems and compared them with simple baselines and traditional NLP models.
- Our empirical evaluation (Section 6) shows the neural networks are more accurate in predicting the query error class, achieving a higher F-measure on classes with fewer samples. For run-time and answer size prediction, the neural networks obtain better results, particularly on complex queries. Additionally, we found character level CNNs are able to generalize better under various problem settings.

Section 7 describes related work and Section 8 concludes.

## **2 MOTIVATING EXAMPLE**

We use the Sloan Digital Sky Survey (SDSS) as a motivating example. SDSS is an astronomy project that provides a digital map of the sky [44, 45, 49, 50]. SDSS data can be queried through SQL. A diverse set of end users, ranging from high school students to astronomers, with varied levels of astronomy and SQL knowledge, use SDSS [49]. To help users compose queries, SDSS provides some resources. These include a step-by-step tutorial for SQL basics and a set of sample SQL queries that can be used as templates. There are also descriptions of costly queries with hints on how to rewrite them to ensure they run faster. For example, users are advised to always start with a "count" query to estimate of the query answer size and prevent long wait times. Figure 1 shows an inefficient query example, given on the SDSS website, which runs the function dbo.fPhotoFlags('BLENDED') for several records in PhotoObj. The users are advised to rewrite it to a query that invokes the function only once.

While these resources are helpful, they are not sufficient. In particular, the step-by-step SQL tutorial is generic and helps inexperienced users learn SQL syntax rather than write meaningful queries. The sample set of queries is small and static compared to the size of the database and the complexity of potential queries. The SDSS schema has 87 tables, 46 views, 467 functions, 21 procedures, and 82 indices. The schema size and complexity makes it hard for users to become familiar enough with the SDSS database to effectively optimize and tune their queries. Ad-hoc hints do not cover all possible optimization opportunities. In this context, real-time query performance estimates like answer size or execution time can increase user productivity and efficiency.

DBAs are another group who interact with SDSS. One of their tasks is to analyze the incoming requests and queries during a session and decide the class of the end user (e.g., human, bot, program) that generated the queries. This is called the session classification problem, where a session is defined as a sequence of interactions between an end user and the system. Session classification allows DBAs to improve the services offered based on usage patterns [50].

Session classification is challenging. First, identifying individual sessions is difficult. This is because users are anonymous and do not necessarily login to the system, their IP addresses can change, and the same IP address can access different SDSS interfaces. In fact, there has been research on automatic session identification as a separate task [30]. We follow [50] and assume that a session is characterized by an ordered sequence of *hits* (i.e., SQL query or web request) from a single IP address, such that the gaps between hits in the sequence is no longer than 30 minutes [44, 50].

The second step in session classification is specifying the correct label for an identified session. These labels can be used to enforce certain policies and optimize system services (e.g., for resource allocation, or to design different interaction modalities based on the usage patterns of different types of sessions [50]). Although the web requests contain a string that describes the browser or program that generated the request, this "agent string" is not reliable. Consequently, SDSS user sessions are labeled using a combination of agent string, IP address, and behavior during session. This procedure does not consider the content or syntactic properties of queries. Therefore, a question that arises is whether the raw query itself can be used for performing the label assignment. This functionality would provide a complementary resource for

SELECT \* FROM PhotoTag WHERE objId=0x112d075f80360018

(a) Error class: success, session class: bot, answer size: 1, and CPU time: 0.015

```
SELECT p.objid,p.ra,p.dec,p.u,p.g,p.r,p.i,p.z

FROM PhotoObj AS p

WHERE type=6 AND p.ra BETWEEN (156.519-0.2) AND (156.519+0.2) AND

p.dec BETWEEN (62.835-0.2) AND (62.835+0.2)

ORDER BY p.objid
```

(b) Error class: success, session class: browser, answer size: 98877, and CPU time: 41.342999

#### Figure 2: Example SQL queries in SDSS workload.

assisting DBAs. Additionally, it helps automate identification of human traffic, which is needed for downstream usability problems, like query recommendation for end users [30].

# **3 PROBLEM DEFINITION**

We use small letters for scalars, capital letters for sets or sequences, bold small letters for vectors, and bold capital letters for matrices.

DEFINITION 1. A query  $Q = (t_1, ..., t_n)$  is a sequence of tokens from a vocabulary V. We consider two sets of vocabularies: a vocabulary that contains characters and a vocabulary that contains words. We define v to denote the size of V and Q to denote the collection of all queries over V.

DEFINITION 2. For a token  $t_i$  in a query Q, we define  $e_i \in \{0, 1\}^{\upsilon}$  as the one-hot encoding of  $t_i$ , i.e., a vector of bits for tokens in V where the bit that corresponds to  $t_i$  is 1 and all the other bits are 0. We use  $\mathbf{x}_i \in \mathbb{R}^d$  to refer to the distributed representation of  $t_i$  in latent space obtained using an embedding matrix  $X \in \mathbb{R}^{d \times \upsilon}$   $(\mathbf{x}_i = X e_i)$ . We define an *n*-gram in Q as a sequence of *n* tokens that appear in Q.

EXAMPLE 1. For the query in Figure 2a,  $Q_1 = ("SELECT", "*", "FROM", ...)$  is a query with a vocabulary of words, and  $Q_2 = (S', E', L', E', ...)$  is a query with a vocabulary of characters. The one-hot encoding of  $t_1 = "SELECT"$  and  $t_2 = "*"$  w.r.t. a vocabulary  $V = \{t_1, t_2, t_3, t_4\}$  is  $e_1 = [1 \ 0 \ 0 \ 0]$  and  $e_2 = [0 \ 1 \ 0 \ 0]$ . Given an embedding matrix  $X = [[3 \ 6 \ 4] \ [9 \ 5 \ 8] \ [4 \ 3 \ 0] \ [6 \ 0 \ 4]]$ ,  $\mathbf{x}_1 = [3 \ 6 \ 4]$  and  $\mathbf{x}_2 = [9 \ 5 \ 8]$  are the distributed representation of  $t_1$  and  $t_2$ , respectively. The sequence ("SELECT", "\*", "FROM") is a 3-gram of words and ('F', 'R', 'O', 'M') is a 4-gram of characters.

We want to model queries and their properties to generate feedback for end users and DBAs. Similar to [5], we assume a-priori access to the database instance, i.e., the tuples and their statistics, is not available. This is commonly the case for users of web services like Google BigQuery. Instead, our approach exploits the rich content of large query workloads: DEFINITION 3. [Query workload] Let  $\mathcal{W} = \{(Q_i, y_i)\}_{i=1}^n$  denote the input query workload, where  $Q_i$  is a query statement and  $y_i$  is a query label. The label is a query property that is obtained by submitting  $Q_i$  to the database.  $\Box$ 

The query statement  $Q_i$  is typically a SQL query and can contain clauses such as SELECT, EXECUTE, CREATE, ALTER, or combinations such as DELETE | UPDATE | INSERT clauses. However, in realistic workloads such as SDSS, the end user can submit *any* query to the system, including a random natural language sentence. So, the query type is not restricted.

The query label  $y_i$  can correspond to different query properties, e.g., answer size or CPU time. We focus on four types of query labels. Our goal is to develop models that can predict these labels — prior to execution. For each query  $Q_i$  the error class  $y_i^e$  is a numeric indicator of whether the query successfully executed or not. The query total CPU time  $y_i^c$ is a real number and represents the query execution time. The query answer size  $y_i^a$  is an integer and represents the number of rows retrieved for the query. The session class  $y_i^s$  is the class of client that generated the query (and its session). Figure 2 shows sample queries from our SDSS query workload, along with their properties.

DEFINITION 4. [Query Facilitation Problems] Given a query workload  $\mathcal{W} = \{(Q_i, y_i)\}_{i=1}^n$  and a new query  $Q_*$ , a query facilitating problem is to predict the label  $y_*$  of  $Q_*$ . We define four query facilitating problems depending on the label: error classification problem, CPU time prediction problem, answer size prediction problem, and session classification problem, where the label corresponds to either the error classes, CPU times, answer sizes, or the session classes.

The underlying assumption in Definition 4 is that  $Q_*$ and W have similar execution conditions, e.g., run over the same database instance. However, this restriction does not hold in many real applications. For example, cloud-based, multi-tenant, and multi-database platforms receive millions of queries from end users based on hundreds of schemas [23]. It is vital to consider such settings and develop models that generalize well to unseen queries. Formally, we define:

DEFINITION 5. [Query Facilitation Problem Settings] Given a query workload W and a new query  $Q_*$ , the problems in Definition 4 can be studied under the following settings:

- (1) *Homogeneous Instance:*  $Q_*$  and W are posed to the same database instance.
- (2) Homogeneous Schema: Q<sub>\*</sub> and W are posed to different database instances with the same schema.
- (3) *Heterogeneous Schema:*  $Q_*$  and W are posed to different databases with different schemas.

In this definition, we assume W and  $Q_*$  are executed in the same DBMS. However, the definition can be extended to

include settings that vary w.r.t. other execution conditions for W and  $Q_*$ , e.g., their SQL version or their DBMSs. Moreover, as the problem setting heterogeneity increases, the prediction problem becomes more challenging. Our empirical study in Section 6.2 confirms that while the prediction error of all models increases with increasing problem heterogeneity, some models can generalize better across settings.

## 4 WORKLOADS AND ANALYSIS

We describe our workloads in Sections 4.1 and 4.2, analyze them in Section 4.3, and summarize the implications of our analysis on model selection and evaluation, in Section 4.4.

# 4.1 SDSS Workload

The SDSS dataset contains logs of queries and requests submitted to SDSS servers. It is described in [44], which we briefly summarize here. For hits, logged data includes the submitted query statement, the version of the database that was queried, the IP address of the computer that generated the hit, the web agent string which specifies the software system that generated the hit, and a time stamp for the hit [44]. In the SDSS schema, hits are recorded in the "SqlLog" and "Weblog" tables, while session information is recorded in the "Session" and "SessionLog" tables. Additional tables record auxiliary information about the hits. The "SqlLog" table contains around 194 million SQL *query log entries* that are grouped into approximately 1.6 million sessions. We extracted the following information from the SDSS dataset:

- The raw **query statement**, extracted from the "SqlStatement.statement" column. This statement can range from a correct SQL statement to random text.
- The **query CPU time** label, extracted from the "SqlLog.busy" column. This value is a real number and represents the query CPU time in seconds [41].
- The **query answer size** label, extracted from the "Sql-Log.rows" column. This value is an integer and represents the number of rows retrieved for the query.
- The **query error class** label, extracted from the "SqlLog.error" column. The three error classes include success (the numeric value 0 means that the query successfully executed), non\_severe error (the numeric value 1), and severe error (the numeric value -1, indicates an invalid query that was rejected by the web portal and was not submitted to the database server).
- The query session class label is extracted through a series of joins on the following tables in the SDSS schema: WebAgentString, AgentStringID, WebAgent, WebLog, SessionLog, Session, and SqlLog (details in [61]). The seven session classes are no\_web\_hit (the session is not established through the Web), unknown (the session is established through the Web but no agent string is reported),

bot (e.g., search engine crawler), admin (administrative service, e.g., performance monitor), program (a user program, e.g., data downloader), and browser (a web browser).

The large size of the SDSS dataset (including 194 million query logs in the SqlLog table) poses a computational challenge in developing machine learning models. In addition, the SDSS dataset has data redundancy [22, 47]. The first type of redundancy is because many sessions can contain thousands of query logs with the same template for their query statements, e.g., bot sessions or administrative sessions typically submit the same query template but with different constants. The second type of redundancy is caused when the same query statement appears in different query logs, with varying values for properties like session class, error class, and answer size. This is because the same statement can be submitted in different sessions, via different access interfaces, and against different versions of the database.

To resolve the redundancy and size issues, we extract a workload by sampling a subset of the SDSS dataset. For the first redundancy issue we randomly sample a SQL query log from each session to ensure a large and diverse subset (the input of our problems is a raw query statement and is independent of other queries in the same session). The result contains 1, 563, 386 query logs. For the second redundancy issue, we group query logs with the same query statement. We found 18.5% of the query statements appear in more than one query log (see [61]). Therefore, we aggregate their metadata labels. In particular, for answer size and CPU time we use the average of these values as the label. For session class, and error class, we use the majority class as the label (with ties broken randomly). Our final query workload contains 618, 053 unique query statements. Details are in [61].

## 4.2 SQLShare Workload

The SQLShare query workload [22] is the result of a multiyear deployment of a database-as-a-service platform, where users upload their data, write queries, and share their results. This workload represents short-term, ad-hoc analytics over user-uploaded datasets. We use the SQLShare workload in our work and extracted the following information:

- The raw **query statement**, extracted from the "Query" column. This may be a syntactically incorrect SQL query.
- The **query CPU time** label, extracted from the "QExec-Time" column. This value is an integer and represents the query CPU time in seconds.

### 4.3 Workload Analysis

*4.3.1 Query Statement Analysis.* We analyze the query statement properties to understand the type of queries posed and their syntactic properties and statistics. Regarding the

query statement types, SELECT statements comprise approximately 96.5% and approximately 98% of statements on SDSS and SQLShare, respectively. The remaining 3.36% (21540) and 2.02% (544) of statements on SDSS and SQLShare, correspond to types such as EXECUTE, CREATE, DROP, UPDATE, ALTER, and various combinations like DELETE | UPDATE | INSERT.

We used the ANTLR parser [42] to generate the Abstract Syntax Trees (AST) of query statements and extract 10 syntactic properties: (1) the number of characters in a query, (2) the number of words in a query (digits are replaced with the <DIGIT> token), (3) the number of function calls, (4) the number of join operators, (5) the number of unique table names in the query, (6) the number of selected columns in the query, (7) the number of predicates (logical conditions, e.g., s.flags\_s=0) used in a query, (8) the number of predicate columns (table names), (9) the query nestedness level, and (10) a nested aggregation indicator that is true when nested queries involve aggregation.

Statistics of the syntactic properties of SDSS statements are shown in Figure 3 (see [61] for SQLShare). Figures 3a and 3b plot the distribution of characters and words for SDSS. The maximum number of characters and words is 7, 795 and 2,975, respectively. Around 30% of the queries have more than 62 characters, and more than 224 words (which are the corresponding distribution means). Figures 3c-3i report key structural metrics such as the number of joins and number of predicates for SDSS. Approximately 5.91% (1.68%) of the queries in SDSS (SQLShare) have at least one join operator, 14.01% (29.74%) of the queries in SDSS (SQLShare) access more than one table, 0.34% (7.88%) of the queries in SDSS (SQLShare) are nested queries, and 0.03% (0.71%) are nested queries with aggregation. Note that the small percentage of nested queries still corresponds to a considerable number of queries (2, 112 for SDSS and 2, 107 for SQLShare).

Our analysis of the syntax of queries in SDSS and SQL-Share shows that these workloads have queries of various complexity w.r.t. the syntactic properties that we studied. Comparing the syntactic properties of the statements in SDSS with those in SQLShare, we observe that while queries are typically longer in SQLShare, the mean number of predicates in the where clause for SDSS is approximately four times that of SQLShare. Although SQLShare queries access more tables on average, SDSS queries perform more joins on average. Finally, SQLShare's queries are more complex in both nestedness and aggregation.

4.3.2 Label Analysis. Figures 4a and 4b show the label distributions of the classification problems for SDSS. As shown in Figure 4a, the error classes are imbalanced; 97.22% of the queries ran without an error (success), while 1.93% had non\_severe errors, and 0.85% had severe errors. Figure 4b shows that session classes are also imbalanced, e.g., program



Figure 3: Structural properties of SDSS query statements.



Figure 4: Label distributions for classification (Figures 4a and 4b) and regression problems (Figures 4c-4e).

and bot comprise 7.93% and 25.98% of the workload, respectively. Note, a simple model that only predicts the majority class (e.g., success in error classification) will achieve a high accuracy. We address this issue in our evaluations by separately calculating the per-class F-measure.

Figures 4c-4e show the label distributions for regression problems. Figure 4c shows SDSS answer size distribution, which ranges from a minimum of -1 (the query did not run due to an error) to a maximum value of 966, 278, 220 tuples. Despite the wide range of values, the data is concentrated around smaller values with a median of 1, i.e., half of the queries either do not run, return no answer, or return only one answer. Figure 4d shows the CPU time distribution on SDSS. The time ranges between 0 and  $10^8$  seconds with the majority of queries taking little CPU time. Figure 4e shows the CPU time distribution in SQLShare ranges approximately between 0 and  $4 * 10^6$  seconds.

#### 4.4 Workload Analysis Implications

4.4.1 Model Selection and Train Loss Functions. Our query facilitation problems in Definition 4 can broadly be categorized as supervised classification and regression problems.

Text classification in NLP is a closely related area. Traditional NLP models, work in two stages: a feature extraction phase, where input features are hand-engineered, and a prediction phase. As shown in Figure 3, queries range in complexity and extracting an adequate set of features can be challenging. Neural network architectures can learn features automatically. They combine the feature extraction and prediction stages in a joint training task, which allows them to develop features and representations for the task [8]. LSTMs are a type of recurrent neural network (RNNs), and are one of the dominant models for text classification. They treat text as sequential inputs and try to preserve long-term dependencies between tokens. However, query statements are long (Figures 3a and 3b), and this property can negatively affect the performance of LSTMs. As an alternative, we assess CNNs, which are feed-forward networks. Rather than preserving long-term dependencies, CNNs automatically identify local patterns (i.e., n-grams) in the input and preserve them in their feature representations. For NLP tasks, CNNs are known to be competitive with several more sophisticated architectures (e.g., LSTMs) and are easier to train and interpret [3, 54].

Moreover, we observe that SQL queries often contain mathematical expressions consisting of numbers and operators. These expressions significantly affect query properties like query answer size or CPU time [13]. It is beneficial to retain relevant information in the representation of queries. However, the set of variable names and digits used in code snippets is unbounded, and there are many rare words. For word-levels models, this leads to the unbounded or open vocabulary problem, which creates practical issues when learning representations in machine learning [32]. We apply the models in Section 5, at both character and word level. For the latter, we replace the digits with a <DIGIT> token to control for the vocabulary size.

Regarding the train loss functions, we observe that the error class and session class labels are imbalanced (Figures 4a and 4b). For some applications, like bot detection, models that perform accurately on certain classes may be required. Typically, application-dependant assumptions are enforced by either re-sampling the data, or using weighted loss functions during training of models. Because our work does not focus on specific applications (e.g., bot detection), we treat all classes equally and use an unweighted cross entropy loss function for training the classification models in Section 5. Our evaluation (Section 6), however, considers this class imbalance where we analyze performance w.r.t. each class.

The regression labels have a wide range of values and are highly skewed, with the majority of queries concentrated around small values (Figures 4c-4e). To prevent the models from being too sensitive to queries with a large label value (outliers), we perform two steps. We apply a logarithmic transformation to the values of these labels  $y'_i$  =  $\ln(y_i + \epsilon - \min(y))$ , where  $y_i$  is the label of query *i*, and *y* is a vector representing the labels (answer size or CPU time) of all queries, and  $y'_i$  is the log-transformed value. When  $y_i = \min(\mathbf{y}), \epsilon > 0$  prevents the input of the ln(.) function from being zero. We set  $\epsilon = 1$  to make the transformation non-negative. We use the log-transformed values of CPU time and answer size as the labels of queries in regression problems. Moreover, to ensure that the regression models are robust to outliers in the data, during training we use the well-known Huber loss [17], which is a hybrid between  $l_2$ -norm for small residuals and  $l_1$ -norm for large residuals.

4.4.2 Model Evaluation. In our work, we want to help end users write queries. This is particularly important in settings where query statements are complex and for users who have little experience. Therefore, we need to assess model performance on complex query statements. However, statement complexity information is not included explicitly in the data. To effectively assess feasibility of complex queries, we must define both a notion of query statement complexity and a proxy measure that captures it.



Figure 5: SDSS structural property correlations.

Similar to [22], we want a proxy metric for complexity that reflects the cognitive effort required to write the query statement. Metrics based on query run time [13] are not adequate in this context. The run time depends on factors like the load of the database or the size of data selected, which are not relevant to the cognitive effort of the user, e.g., a simple query that selects all rows of a large table can have a long running time. In [22], query complexity is defined in terms of the query's ASCII length and the number of distinct operators in the query execution plan. However, we do not assume access to the database instance or execution plan. The ASCII length, on the other hand, might not be a sufficient proxy for complexity, e.g., when a query has a simple structure with similar operations repeated many times. Additional syntactic properties may be required.

However, it is not clear which set of syntactic properties capture a meaningful notion of statement complexity. Figure 5 shows the correlation matrix for the syntactic properties in Figure 3 for SDSS. We observe some properties are positively and linearly correlated with other types of properties, and hence indicative of them. For example, the number of characters is linearly correlated with the number of words, the number of predicate columns, and the number of predicates. So the latter properties are redundant since the number of characters is indicative of them. But number of characters is not positively correlated with properties like nested aggregation and nestedness level and may not capture those complexities. As another example, the number of joins is only linearly correlated with the number of unique table names. We observed similar patterns for SQLShare (see [61]).

Overall, these observations suggest that a subset of syntactic properties might be required to capture the full range of potential query complexities. Based on the query statement feature correlations shown in Figure 5, we chose a subset of five syntactic properties for the qualitative analysis in



Figure 6: SDSS analysis by session class. The top and bottom of each box represent the first and third quartiles of the class data distribution. The horizontal line in each box is the median, the mean is a green triangle.

Section 6.3.3. These include the number of characters, the number of functions, the number of joins, the nestedness level, and the nested aggregation indicator.

One viable assumption may be that different classes of users write queries with different complexity levels. For example, bot queries may use linear predicates in the where clause, while queries via browser may be more complex. Thus, session class, if available, can be an indirect proxy for query complexity. Figure 6a plots the distribution of the answer size for each session class. The no\_web\_hit and browser classes have similar distributions, with the latter having slightly smaller values (likely due to the limitations for queries posed via the web-based interface). In Figure 6b the distribution of the CPU time for each session class is shown. Queries in the no\_web\_hit class have a wider range of values. Figures 6d and 6c show the query size distributions by session class. Overall, queries from no\_web\_hit and browser classes have similar distributions at both the character and word level. These figures suggest that queries in the no\_web\_hit and browser class are more complex.

# **5 METHODS**

Based on our analysis in Section 4.4.1, we extend three models from the NLP domain and benchmark their performance for our problems. In Section 5.1 we describe a traditional model. In Section 5.2 we describe a three-layer LSTM model and in Section 5.3 we describe a shallow CNN.

# 5.1 Traditional Model

Traditional machine learning models work in two stages: a feature extraction phase and a prediction phase. For the feature extraction phase, Bag-of-ngrams and its TFIDF (termfrequency inverse-document-frequency) are commonly used in NLP applications. For the Bag-of-ngrams, we select the most frequent *n*-grams (up to 5-grams) from the training set. These features comprise the domain vocabulary V with size v. Thus, this representation maps each query to a vdimensional vector obtained by computing the sum of the one-hot representation of the *n*-grams that appear in the query. Next, we compute the TFIDF weight of each token  $t_i$  in the *v*-dimensional representation of query *Q* w.r.t. the collection of queries Q. In particular, the weight of token  $t_i$  is computed using  $TFIDF(t_i, Q, Q) = TF(t_i, Q) \times IDF(t_i, Q)$ . Here  $TF(t_i, Q)$  is the normalized frequency of  $t_i$  in Q. The normalization prevents bias towards longer queries. The  $IDF(t_i, Q)$  component describes the discriminative power of  $t_i$  in Q and helps to control for the fact that some tokens are generally more common than other tokens. It can be computed by  $\log \frac{|Q|}{1+|\{Q \in Q, t_i \in Q\}|}$ , where the denominator is the number of queries in Q that contain  $t_i$ . The TFIDF value increases proportionally to the frequency of a token in a query but is counterbalanced by the frequency of the term in the collection.

We then apply a prediction model given this fixed vdimensional feature vector. For classification problems, we apply the multinomial logistic regression model. For regression problems, we use Huber loss [17]. We implemented this model using scikit-learn [43].

# 5.2 Three-layer LSTM

Long-Short Term Memory (LSTM) is a type of recurrent neural network (RNN) [57]. RNNs can process sequential inputs of arbitrary length. Standard RNN units work by reading the input sequence one token at a time from left to right. At every step *i*, a hidden state  $h_i \in \mathbb{R}^k$  is emitted, which is a semantic representation of the sequence of tokens observed so far. Specifically,  $h_i$  is produced using the recurrent equation  $h_i = f(Wx_i + Uh_{i-1} + b)$  where  $x_i \in \mathbb{R}^d$  is the distributed representation of the input token  $q_i$ , and  $h_{i-1} \in \mathbb{R}^k$  is the hidden state at i - 1. The parameters of this RNN unit include weight matrices W and U, and a bias vector b. f(.) is a point-wise non-linear activation function, such as the Sigmoid or Rectified Linear unit (Relu) [14].

Standard RNNs suffer from the *vanishing gradient* problem. In particular, during training, the gradient vector can grow or decay exponentially [14, 51]. LSTMs are a more effective variant of RNNs. They are equipped with a memory cell  $c \in \mathbb{R}^k$ , which helps preserve the long-term dependencies better than standard RNNs. The LSTM unit [57] has a hidden state  $h_i$  that is a partial view of the unit's memory cell. The unit is equipped with additional parameters and machinery to produce  $h_i$  from  $c_{i-1}$  (memory cell at step *i*),  $x_i$ , and  $h_{i-1}$ (details in [61]).

Since well-known RNN architectures do not exist for our problems, we explore those used in similar domains. Deep architectures consisting of many layers are often developed to learn hierarchical representations for the input and to learn non-linear functions of the input [8]. However, increasing the number of layers and units increases the number of parameters to learn, and training time increases substantially.

A two-layer character-level LSTM architecture was used to predict program execution in [57]. We found a three-layered LSTM model performed better for our problems. We use the output of the last layer as the query vector representation. For classification problems, we apply the softmax operation to generate the output probability distribution. Similar to the traditional models, we use the cross-entropy loss for classification problems. For regression problems, we pass the vector through a linear unit and use Huber loss. To optimize the network, we examined both Adam and AdaMax [33] which are gradient-based optimization techniques that are well suited for problems with large data and many parameters. We found the latter performed better.

#### 5.3 Shallow CNN

Convolutional Neural Networks (CNNs) are feed-forward neural networks that process data with grid-like topology, e.g., a sequence of concatenated distributed representations of tokens in NLP. Their application in NLP enables the model to extract the most important n-gram features from the input and create a semantic representation. As a result, long-term dependencies may not be preserved and token order information is preserved locally. CNNs, however, have comparable performance to RNNs, they are easier to train, and are also parallelizable [3, 54, 56].

$$\begin{bmatrix} 0 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \\ x_{1:n} & & & \\ \end{bmatrix} * \begin{bmatrix} 1 & 0 & 1 \\ w & & \\ x_{1:n} & * & \\ \end{bmatrix}$$

Figure 7: 1D convolution operation with input  $x_{1:n}$  and kernel *w*. The output  $p = x_{1:n} * w$  is produced by sliding *w* over  $x_{1:n}$  and computing the dot product.

Each layer in a CNN consists of three stages [14]: a convolution stage, a detection stage, and a pooling stage. We explain each of stages based on a 1D convolution operation, although higher dimensions are also possible.



Figure 8: 1-layer CNN, adapted from [31].

The convolution stage applies several convolution operations. A convolution operator has two operands: a multidimensional array of weights, called the kernel, and a multidimensional array of input data. Convolving the input with the kernel consists of sliding the kernel over all possible windows of the input. At every position *j*, a linear activation  $p_i$  is obtained by computing the dot product between the kernel entries and local regions of the input. Figure 7 shows an example of a 1D convolution operation. Let  $\oplus$  denote the concatenation operation, and  $\mathbf{x}_{1:n} = \mathbf{x}_1 \oplus \mathbf{x}_2 \oplus \cdots \oplus \mathbf{x}_n$  denote the concatenated distributed representations  $x_i$  in an input query (stacked length-wise as a long column). Let  $x_{i:i+m-1}$ represent a window of *m* words and  $w \in \mathbb{R}^m$  denote a kernel. The dot product of *w* and each *m*-gram in the sequence is computed to obtain  $p_i = \mathbf{w}^T \mathbf{x}_{i:i+m-1} + b$  where  $b \in \mathbb{R}$  is a bias term. By sliding the kernel over all possible windows of the input, we obtain a sequence  $p \in \mathbb{R}^{n-m+1}$  (b = 0 in Figure 7). Note, in the convolution stage, several kernels with varying window sizes may be convolved with the data to produce different linear sequences. In the detector stage, the linear sequence *p* is run through a non-linear activation function. This results in a sequence of non-linear activations called the activation map a = f(p), where  $a \in \mathbb{R}^{n-m+1}$  and f(.) is a non-linear activation function, e.g., Relu. In the pooling stage a pooling operation is applied to the activation map to summarize its values which also enables the model to handle inputs of varying size. For example, the max-pooling function returns the maximum, i.e.,  $q = \max\{a\}$ .

Figure 8 shows the shallow CNN architecture in [31], which we adapt for our application. The input query  $\mathbf{x}_{1:n}$  and convolution operations are shown in 2D for easier presentation. In the convolution stage, several filters of varying window size  $m \in \{3, 4, 5\}$  are applied, and the resulting sequence  $\mathbf{p}_1, \mathbf{p}_2, \ldots$  is passed through a Relu function to generate activations  $\mathbf{a}_1, \mathbf{a}_2, \ldots$ . Different size sequence inputs and kernels result in activation maps of different sizes. To deal with variable length of the input and also obtain the most important feature, a max pooling operation is applied to obtain a single feature per kernel,  $q = \max{a}$ . The resulting

features for all kernels are concatenated to produce a fixed size vector  $\boldsymbol{g} \in \mathbb{R}^{K}$ , where *K* is the number of kernels. This output is used to create a fully connected layer, followed by a dropout layer.

We tried changing the architecture by increasing the number of kernels and the window size but did not obtain significant improvements. Similar to the other models, for classification problems we apply the softmax operation to generate the output probability distribution. We use the cross-entropy loss for classification problems. For regression problems, we pass the vector through a linear unit and use Huber loss to learn the parameters. We used AdaMax as the optimizer.

# 6 EMPIRICAL EVALUATION

We evaluate the performance of models in Section 5 on the four query facilitation problems, considering two aspects: (1) the different query facilitation problem settings and (2) the query statement complexity (described in Section 4.4.2).

## 6.1 Setup

**Data split.** For Homogeneous Instance, we used our extracted SDSS workload. For Homogeneous Schema, we used SQLShare. In both settings, we randomly split the queries. For Heterogeneous Schema, we used SQLShare and randomly split the data based on users, so as to decrease the likelihood of data sharing. The SDSS workload has a total of 618, 053 queries. SQLShare has a total of 26, 728 queries. In all three settings, we used a (80/10/10) split for the train, validation, and test sets, respectively.

**Methods compared.** We compare the models in Section 5, where character-level model names begin with c and word-levels models with w. The traditional models are ctfidf and wtfidf. The 3-layer LSTM models are clstm and wlstm. The CNN models are ccnn and wcnn. For each prediction task, we also include a simple baseline. For classification problems, mfreq predicts the most frequent label, i.e., it predicts success for error classification, and no\_web\_hit for session class prediction. For regression problems, median predicts the median of the corresponding train distribution, i.e., the median of the train CPU time distribution is 0. Following [1, 13, 38] we report results for an opt model which uses linear regression to predict CPU time from the query optimizer estimates cost estimates.

**Hyper-parameter tuning.** We tune the hyper-parameters based on Homogeneous Instance (SDSS). To keep this problem tractable, we restrict the set of hyper-parameters of each model and choose the best set of hyper-parameters based on performance on the validation set. We fixed the learning rate 1e-3, batch size to 16, and token embedding size to 100. For

clstm and wlstm, we tested the number of hidden dimensions in {150, 300} and clipping rate in {0.25, 0}. For ccnn and wcnn, we tested number of kernels in {100, 250}, drop out in {0.5, 0}, and clipping rate in {0.25, 0}. We report results for the best performing model and also use the model in Homogeneous Schema and Heterogeneous Schema settings.

Performance metrics. For the classification problems we report the test average loss measured by cross entropy loss . We also report accuracy, which is the number of correct predictions divided by the total number of predictions. Due to the class imbalance for both error and session classification, for every class C, we report the per class F-measure computed by  $F_C = \frac{2.\operatorname{Precision}_C.\operatorname{Recall}_C}{\operatorname{Precision}_C+\operatorname{Recall}_C}$ . Precision<sub>C</sub> is the number of correct predictions for class C divided by the total number predictions for class C. Recall<sub>C</sub> is the number of correct predictions for class C divided by the total number queries in class C. For the regression problems we report the test average loss measured by Huber loss. We also report Mean Square Error (MSE) computed as MSE =  $\frac{1}{m} \sum_{i}^{m} (y'_{i} - \hat{y}_{i})^{2}$ , where  $y'_i$  is the log-transformed label (CPU time or answer size) of query *i*, and  $\hat{y}_i$  is the predicted value. We also use qerror [36] to measure the quality of estimates. It is defined as qerror<sub>i</sub> = max $(\frac{y_i}{\hat{y}_i}, \frac{\hat{y}_i}{y_i})$  for a query. For qerror analysis, we do not apply log-transformation to the query labels.

#### 6.2 Model Performance

6.2.1 Homogeneous Instance. Table 1 (left) shows the error classification results. The mfreq baseline achieves a high F<sub>success</sub> but performs poorly w.r.t. other classes. All other models improve upon this baseline. The ccnn model obtains a high  $F_{severe} = 0.7961$  and has the highest test accuracy. Table 1 (middle) shows results for CPU time prediction. The wcnn model obtains the lowest test loss, followed closely by the wlstm model. Table 1 (right) shows results for answer size prediction. The ccnn model obtains the lowest test loss, followed closely by the clstm model. In Section 6.3.1 we assess both problems w.r.t. MSE values. Table 3 shows results for session classification. Again the mfreq baseline achieves a high Fno web hit (the majority class) but under-performs all other models w.r.t. other classes. The highest test accuracy is obtained by the ctfidf model, which outperforms other models in the F-measure of individual classes, except for Funknown which is 0. Fadmin is 0 since admin only has 2 queries in the test set.

Table 2 shows the qerror of the answer size predictions of 62K test queries in SDSS. The percentage of queries that have at most the reported qerror is shown, e.g., the qerror of 75% of the test queries is less than 2.38 for clstm. Intuitively, qerror for answer size is the factor by which an estimate differs from the true answer size. We observe ccnn and clstm have the lowest gerrors. Note, all models perform well for 50% of

Error Classification							CPU Time		Answer Size		
Model	υ	p	Accuracy	Fsevere	F <sub>success</sub>	$\mathbf{F}_{\text{non}\_severe}$	Loss	p	Loss	p	Loss
baseline	-	-	0.9730	0.0000	0.9863	0.0000	0.5951	-	0.0675	-	1.6357
ctfidf	500000	1500000	0.9778	0.7131	0.9888	0.0053	0.5860	500000	0.0668	500000	1.0400
ccnn	159	17403	0.9797	0.7961	0.9897	0.1669	0.1106	16801	0.0471	16801	0.7517
clstm	159	1944003	0.9786	0.6922	0.9893	0.2206	0.0830	1943401	0.0452	529651	0.7678
wtfidf	500000	1500000	0.9773	0.6546	0.9885	0.0620	0.5836	500000	0.0668	500000	1.0922
wcnn	85942	8597953	0.9790	0.7441	0.9894	0.2006	0.1006	8595101	0.0441	8595101	0.8472
wlstm	85942	10522303	0.9776	0.6971	0.9887	0.0018	0.0691	10521701	0.0443	9107951	0.8256

Table 1: Query error classification (left), CPU time (middle), and answer size (right) prediction in Homogeneous Instance (SDSS). Here v is #tokens in the vocabulary, p is #model parameters, Loss is the average test loss (lower loss is better). baseline is median in the regression problems and mfreq in the classification problem. In the classification problem,  $F_C$  is the F-measure of class C. The #samples of each class in the test set are: severe = 533, success = 60138, non\_severe = 1134.

Model	50%	75%	80%	85%	90%	95%
median	1	36	50	144	1885	50000
ctfidf	1.13	4.86	10	25	88	727
ccnn	1.36	2.60	3.75	6.79	18	174
clstm	1.07	2.38	3.50	6.79	19	172
wtfidf	1.00	5.37	11.04	31.98	100	879
wcnn	1.33	3.42	5.14	10.93	36	295
wlstm	1.12	2.62	4.27	10.43	30	292

Table 2: Answer size prediction qerror (SDSS).

the queries and the main comparison is for the other 50% of the queries for which prediction is more difficult.

6.2.2 Homogeneous Schema. Table 4 reports performance for CPU time prediction in Homogeneous Schema. The ccnn model outperforms other models. Compared to Homogeneous Instance, the overall loss value is higher for all models. This is because the latter poses an additional challenge where the distribution of the queries in individual database instances is different, and to get accuracy compared to Homogeneous Instance, we need to increase model capacity (e.g., add more layers in the architecture). Moreover, observe that the opt model, that is based on the query optimizer cost model, is closer to median in it's error. Our qerror analysis for 2,674 test queries in SQLShare shows ccnn performs better across different percentiles. For 50% and 75% of the queries, qerror is less than 1.94 and 27, resp (see [61]).

6.2.3 *Heterogeneous Schema*. Table 4 reports performance for CPU time prediction in Heterogeneous Schema. Similar to Homogeneous Schema, the ccnn model outperforms others. However, compared to Homogeneous Instance and Homogeneous Schema, the loss value achieved by all models is higher. This is expected since the data is extracted from databases with different schemata, which makes it more challenging for the models to predict, i.e., the train and test sample distributions are different. For opt, prediction is more difficult, too. As explained in [1], the query optimizer cost model assumes I/O is most time consuming, even though certain computations (e.g., nested aggregates over numeric types) are performed in memory. Moreover, a non-linear regression model may improve performance of opt. Our qerror analysis (see [61]) shows ccnn performs better across different percentiles in Heterogeneous Schema. For 30% of the queries, qerror is less than 34. The substantial qerror increase means prediction is harder in Heterogeneous Schema.

6.2.4 Discussion. We found the following: (1) Character-level models (ccnn and ctfidf) obtain the best performance for all problems except CPU time prediction in Homogeneous Instance, where word-level models (wcnn and wlstm) obtained the lowest test loss and MSE. Intuitively, as the problem heterogeneity increases, the number of rare words increases, making it difficult to learn word-level patterns. In Homogeneous Instance setting, however, queries have more words in common (e.g., table names and SQL keywords), and the models can learn the underlying distributions better. (2) Overall, CNN and LSTM architectures outperform others on all problems except session classification, where ctfidf obtains better results in predicting several classes. The frequency of the classes (see Table 3) shows that ctfidf performs better for majority classes (e.g., no\_web\_hit and bot); and CNN and LSTM beat ctfidf in non-frequent classes (e.g., unknown and program) where prediction is more difficult. In addition, ccnn achieves almost the same overall accuracy with much fewer parameters (16801 vs 500000). The neural networks learn features w.r.t. task, but ctfidf and wtfidf are limited to pre-determined features. (3) Regarding generalization of a single model under various

Model	υ	р	Loss	$F_{\text{no}\_web\_hit}$	$\mathbf{F}_{\text{unknown}}$	$\mathbf{F}_{\texttt{bot}}$	$\mathbf{F}_{\text{program}}$	Fanonymous	$\mathbf{F}_{\text{browser}}$	Accuracy
mfreq	-	-	1.7848	0.6186	0.0000	0.0000	0.0000	0.0000	0.0000	0.4478
ctfidf	500000	3500000	1.5786	0.6235	0.0000	0.7272	0.6128	0.6176	0.5618	0.6421
ccnn	159	18607	0.7960	0.5921	0.2373	0.6940	0.6076	0.5441	0.5389	0.6152
clstm	159	1945207	0.8600	0.5774	0.0455	0.6817	0.6366	0.4868	0.5451	0.6102
wtfidf	500000	3500000	1.6077	0.5487	0.2903	0.6958	0.6241	0.6042	0.5416	0.6068
wcnn	85942	8596907	0.8373	0.5411	0.3778	0.6955	0.5344	0.5970	0.5615	0.6004
wlstm	85942	10523507	0.8452	0.5558	0.0000	0.6628	0.6278	0.4482	0.5303	0.5911

Table 3: Query Session classification in Homogeneous Instance (SDSS). The #examples in the test set for each class is: no\_web\_hit = 27677, unknown = 42, bot = 16148, program = 4882, anonymous = 467, browser = 12587.

		Homoge Sche	eneous ma	Heterogeneous Schema		
Model	υ	p	Loss	p	Loss	
median	-	-	2.0049	-	2.1616	
opt	-	-	1.8909	-	2.2841	
ctfidf	500000	500000	0.4742	500000	1.2360	
ccnn	101	11001	0.4625	10901	1.1547	
clstm	101	1937601	0.7935	1937501	1.6046	
wtfidf	500000	500000	0.4898	500000	1.9702	
wcnn	14843	1485201	0.5230	1395901	2.0416	
wlstm	14843	3411801	0.8081	3322501	1.8546	

Table 4: Query CPU time prediction (SQLShare)

settings, ccnn identifies local sequential character patterns which help it learn the underlying data distribution better.

## 6.3 Detailed Qualitative Analysis

6.3.1 Performance by Session Class. We refine the analysis and use session class as a proxy for complexity in Homogeneous Instance. We analyze answer size prediction in Figure 9a, which shows MSE by session class. The MSE trends show that predicting answer size for no\_web\_hit, program, and browser is more difficult. Moreover, the simple baseline median under-performs all models across all sessions. We observe ctfidf and wtfidf also under-perform other models (we found they perform similarly to median for CPU time prediction, see [61]). Our results indicate the neural network models perform better on complex session classes.

6.3.2 Performance by Structural Properties. Figures 9b-9e examine the error of answer size prediction for varying structural properties in Homogeneous Instance. As expected, error increases for more complex queries (with larger number of characters, number of functions, and number of joins). The decrease of error in the middle and end of the graphs in Figures 9b-9d is due to fewer answers for the corresponding queries, which makes prediction easier for all the models (including median, which supports this claim). Figure 10 examines the error of CPU time predication. Figures 10a, 10c, and 10e show the MSE of all models increases as the data heterogeneity (Homogeneous Instance to Homogeneous Schema to Heterogeneous Schema) and the query complexity, i.e., the number of characters, increase. Figures 10b, 10d, and 10f show a breakdown of the prediction error of ccnn, w.r.t. nestedness level, under various settings. The unexpected decrease in MSE of queries with high nestedness (nestedness level=3, 4) is due to better prediction for the few queries with small CPU time.

6.3.3 Case Study. We study performance for two sample queries with different structural properties.  $Q_1$  in Figure 11 is a large query (number of characters =1, 247 and number of words =376) that joins three large tables (e.g., Specobj and Photoobj contain 4, 311, 571 and 794, 328, 715 rows, respectively), selects 49 columns in the answer, and calls 3 functions. The query is from the browser and ran successfully (error class: success) with CPU time of 105.37 sec and returned 304 answers. Comparing ccnn and clstm, the former predicts 116.40 sec for CPU time while clstm's estimation is 980 sec. The query length makes it hard for clstm to capture the long-term dependencies, whereas ccnn detects local patterns and combines them globally to make a better prediction.

 $Q_2$  in Figure 12, is shorter than  $Q_1$  (number of characters =645 and number of words =181), but it is more complex (nestedness level=3, number of functions=5, and number of predicates =11). The query runs instantly since it accesses tables (Jobs, Users, Status and Servers) with fewer rows. Its answer size is 27 rows. The CPU time prediction of ccnn, wcnn and clstm are 1.00 sec, 1.28 sec and 1.01 sec, respectively. Their answer size predictions are 45, 46 and 49. For  $Q_2$  all 3 models perform fairly well. The small CPU time and the answer size of  $Q_2$  compared to  $Q_1$ , contributes to more accurate predictions (due to the logarithmic label transformation and Huber loss (cf. Section 4.4.1)).  $Q_2$  is shorter in length compared to  $Q_1$ , which makes it easier for clstm to make predictions.



Figure 9: Error analysis of answer size prediction in Homogeneous Instance (SDSS). Performance breakdown across session classes in Figure 9a (legend shows the total MSE for each model) and w.r.t. different structural properties in Figures 9b-9d. Error analysis for ccnn by nested aggregation in Figure 9e.

6.3.4 Discussion. Using session information, CPU time and answer size prediction were more difficult for no\_web\_hit, program, and browser sessions, for all models. This is because queries in these classes are more complex compared to other classes (see Figure 6) and are likely issued by humans. Our evaluation by structural properties showed that predicting labels is more difficult for complex queries (e.g., with large number of characters, number of functions, and number of joins), and in settings where data is from heterogeneous sources. Word-level models suffer from many rare tokens in heterogeneous settings, and do not generalize well. Among the character-level models, clstm is sensitive to the query



Figure 10: Error analysis of CPU time prediction by number of characters (left). Error analysis of CPU time prediction for ccnn by nestedness level (right).



#### Figure 11: Sample query Q<sub>1</sub>



Figure 12: Sample query Q<sub>2</sub>

length and is outperformed by ccnn as statement complexity increases.

# 7 RELATED WORK

Deep learning, Machine Learning, and NLP. RNNs and CNNs are dominant in many text applications [8]. Characterlevel LSTMs were used for program execution in [57]. In [31], a one-layer word-level CNN model outperformed tree-structured models that use syntactic parse trees as their input, for text categorization. Deep character-level CNN models [8, 27, 59] outperformed shallow word-level CNNs [27]. Although shallow word-level models have more parameters and need more storage, their computations are faster. Subsequently, deep word-level CNNs have been applied in [29]. LSTMs and CNNs are compared in [3, 54]. CNNs outperform RNNs, for sequence modeling tasks. They are also parallelizable, which leads to speeds up in their execution [54]. We examined both character-level and word-level LSTMs and CNNs. Our work is related to machine learning for Big Code and naturalness [2]; we leave that analysis for future work.

**Deep learning in databases.** Research problems at the intersection of deep learning and databases are introduced in [53]. Examples include query optimization and natural language query interfaces [53]. A feed-forward neural network (with 1 hidden layer) for cardinality estimation of simple range queries (without joins) is proposed and evaluated on a synthetic dataset [39]. Recently, [60] developed a natural language interface for database systems using deep neural networks. In [21], an LSTM autoencoder and a paragraph2vec model were applied for the tasks of query workload summarization and error prediction, with experiments on Snowflake, a private query workload, and TPC-H [52]. Compared to the datasets in [21, 60], SDSS and SQLShare are publicly-available and real-world.

**Modeling SQL query performance.** Estimates of SQL query properties and performance are used in admission control, scheduling, and costing during query optimization. Commonly, these estimates are based on manually constructed cost models in the query optimizer, which may be imprecise. Prior work has used machine learning to accurately estimate SQL query properties [1, 13, 36, 38, 39]. Most works use relatively small synthetic workloads, like TPC-H and TPC-DS, along with traditional two-stage machine learning models. Their results are better with query execution plans as input. Similar to us, the database-agnostic approach in [23] automatically learns features from large query workloads rather than devising task-specific heuristics and feature engineering for pre-determined conditions. However, they focus on index selection and security audits.

**Facilitating SQL query composition.** Form-based [24] and keyword queries [5, 58] can help users write queries. Natural language interfaces, like NaLIR [37], allow complex

query intents to be expressed. Query recommendation by mining query logs [6, 11, 30] is another approach. Given the schema, tuples, and some keywords, the approach in [12] suggests SQL queries from templates. Additional query results are recommended for each query in [48]. However, other than [30], these works access tuples. Other work assume the user is familiar with samples in the query answer. AIDE [9] helps the user refine linear queries using decision tree classifiers. Finding minimal project join queries based on a sample table of tuples contained in the query answer, is studied in [46]. These works are complementary to ours.

**Mining SQL query workloads.** Several usability works use the TPC-H benchmark dataset [52]. TPC-H has 8 tables, contains (22) ad-hoc queries, and data content modifications. A synthetic workload can be simulated from the ad-hoc queries. Query workloads are also used for tasks like index selection [21], improving query optimization [38], and workload compression [7]. Workload compression techniques can provide an orthogonal extension for data extraction part of our work. SDSS has been used to identify user interests and access areas within the data space [40]. Ettu [34], is a system that identifies insider attacks, by clustering SQL queries in a query workload. We focus on different problems.

# 8 CONCLUSION

We address facilitating user interaction with the database by providing insights about SQL queries — prior to query execution. We leverage (only) the abundant information in large-scale query workloads. We conduct an empirical study on SDSS and SQLShare query workloads and adapt various data-driven machine learning models. We found the neural networks (character-level CNNs in particular) outperformed other models, for query error classification, answer size prediction, and CPU time prediction.

There are several avenues for future work. We intend to apply transfer-learning ideas to improve ccnn under heterogeneous settings [4, 55]. More sophisticated models, e.g., deep character CNNs [8] or tree-structured architectures [51] may lead to performance gains. Query workload extraction is another direction. The SDSS dataset is large and noisy. To understand the challenges, we extracted a sample and analyzed our problems. However, more adequate query workloads can be extracted, separately, for various problems. Another direction is to use multi-task models that learn correlations between the query labels, although our models are applicable in broader settings where workloads have only one label. Incorporating other types of meta-data, e.g., the database version that was queried, may increase accuracy. We leave addressing these challenges for future work.

#### REFERENCES

- Mert Akdere, Ugur Çetintemel, Matteo Riondato, Eli Upfal, and Stanley B Zdonik. 2012. Learning-Based Query Performance Modeling and Prediction. In *ICDE*. 390–401.
- [2] Miltiadis Allamanis, Earl T Barr, Premkumar Devanbu, and Charles Sutton. 2018. A Survey of Machine Learning for Big Code and Naturalness. *Comput. Surveys* 51, 4 (2018), 81.
- [3] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. 2018. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. ArXiv 1803.01271 (2018).
- [4] Yoshua Bengio. 2011. Deep Learning of Representations for Unsupervised and Transfer Learning. In UTLW. 17–37.
- [5] Sonia Bergamaschi, Elton Domnori, Francesco Guerra, Raquel Trillo Lado, and Yannis Velegrakis. 2011. Keyword Search over Relational Databases: A Metadata Approach. In SIGMOD. 565–576.
- [6] Gloria Chatzopoulou, Magdalini Eirinaki, and Neoklis Polyzotis. 2009. Query Recommendations for Interactive Database Exploration. In SSDBM. 3–18.
- [7] Surajit Chaudhuri, Ashish Kumar Gupta, and Vivek Narasayya. 2002. Compressing SQL Workloads. In SIGMOD. 488–499.
- [8] Alexis Conneau, Holger Schwenk, Loïc Barrault, and Yann Lecun. 2016. Very Deep Convolutional Networks for Text Classification. ArXiv 1606.01781 (2016).
- [9] Kyriaki Dimitriadou, Olga Papaemmanouil, and Yanlei Diao. 2014. Explore-by-Example: An Automatic Query Steering Framework for Interactive Data Exploration. In SIGMOD. 517–528.
- [10] Bailu Ding, Sudipto Das, Ryan Marcus, Wentao Wu, Surajit Chaudhuri, and Vivek Narasayya. 2019. AI Meets AI: Leveraging Query Executions to Improve Index Recommendations. In SIGMOD. 1241–1258.
- [11] Magdalini Eirinaki, Suju Abraham, Neoklis Polyzotis, and Naushin Shaikh. 2014. QueRIE: Collaborative Database Exploration. *TKDE* 26, 7 (2014), 1778–1790.
- [12] Ju Fan, Guoliang Li, and Lizhu Zhou. 2011. Interactive SQL Query Suggestion: Making Databases User-Friendly. In *ICDE*. 351–362.
- [13] Archana Ganapathi, Harumi Kuno, Umeshwar Dayal, Janet L Wiener, Armando Fox, Michael Jordan, and David Patterson. 2009. Predicting Multiple Metrics for Queries: Better Decisions Enabled by Machine Learning. In *ICDE*. 592–603.
- [14] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press.
- [15] Google. 2010. Big Query. https://cloud.google.com/bigquery/.
- [16] Hakan Hacigümüş, Bala Iyer, Chen Li, and Sharad Mehrotra. 2002. Executing SQL over Encrypted Data in the Database-Service-Provider Model. In *SIGMOD*. 216–227.
- [17] Peter J Huber et al. 1964. Robust estimation of a location parameter. *The Annals of Mathematical Statistics* 35, 1 (1964), 73–101.
- [18] Amazon Inc. 2012. Amazon Redshift. https://aws.amazon.com/ redshift/.
- [19] Snowflake Inc. 2012. Snowflake. https://www.snowflake.com/.
- [20] HV Jagadish, Adriane Chapman, Aaron Elkiss, Magesh Jayapandian, Yunyao Li, Arnab Nandi, and Cong Yu. 2007. Making Database Systems Usable. In SIGMOD. 13–24.
- [21] Shrainik Jain and Bill Howe. 2018. Query2Vec: NLP Meets Databases for Generalized Workload Analytics. ArXiv 1801.05613 (2018).
- [22] Shrainik Jain, Dominik Moritz, Daniel Halperin, Bill Howe, and Ed Lazowska. 2016. SQLShare: Results from a Multi-Year SQL-as-a-Service Experiment. In SIGMOD. 281–293.
- [23] Shrainik Jain, Jiaqi Yan, Thierry Cruane, and Bill Howe. 2018. Database-Agnostic Workload Management. ArXiv 1808.08355 (2018).
- [24] Magesh Jayapandian and HV Jagadish. 2009. Automating the Design and Construction of Query Forms. *TKDE* 21, 10 (2009), 1389–1402.

- [25] Rie Johnson and Tong Zhang. 2014. Effective Use of Word Order for Text Categorization with Convolutional Neural Networks. ArXiv 1412.1058 (2014).
- [26] Rie Johnson and Tong Zhang. 2015. Semi-Supervised Convolutional Neural Networks for Text Categorization via Region Embedding. In Advances in Neural Information Processing Systems. 919–927.
- [27] Rie Johnson and Tong Zhang. 2016. Convolutional neural networks for text categorization: Shallow Word-Level vs. Deep Character-Level. *ArXiv* 1609.00718 (2016).
- [28] Rie Johnson and Tong Zhang. 2016. Supervised and Semi-Supervised Text Categorization using LSTM for Region Embeddings. In *ICML*. 526–534.
- [29] Rie Johnson and Tong Zhang. 2017. Deep Pyramid Convolutional Neural Networks for Text Categorization. In ACL. 562–570.
- [30] Nodira Khoussainova, YongChul Kwon, Magdalena Balazinska, and Dan Suciu. 2010. SnipSuggest: Context-Aware Autocompletion for SQL. PVLDB 4, 1 (2010), 22–33.
- [31] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. ArXiv 1408.5882 (2014).
- [32] Yoon Kim, Yacine Jernite, David Sontag, and Alexander M Rush. 2016. Character-Aware Neural Language Models.. In AAAI. 2741–2749.
- [33] Diederik P Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. ArXiv 1412.6980 (2014).
- [34] Gokhan Kul, Duc Luong, Ting Xie, Patrick Coonan, Varun Chandola, Oliver Kennedy, and Shambhu Upadhyaya. 2016. Ettu: Analyzing Query Intents in Corporate Databases. In WWW. 463–466.
- [35] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. 1989. Backpropagation Applied to Handwritten Zip Code Recognition. *Neural computation* 1, 4 (1989), 541–551.
- [36] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter Boncz, Alfons Kemper, and Thomas Neumann. 2015. How Good Are Query Optimizers, Really? *PVLDB* 9, 3 (2015), 204–215.
- [37] Fei Li and HV Jagadish. 2014. Constructing an Interactive Natural Language Interface for Relational Databases. PVLDB 8, 1 (2014), 73–84.
- [38] Jiexing Li, Arnd Christian König, Vivek Narasayya, and Surajit Chaudhuri. 2012. Robust Estimation of Resource Consumption for SQL Queries using Statistical Techniques. *PVLDB* 5, 11 (2012), 1555–1566.
- [39] Henry Liu, Mingbin Xu, Ziting Yu, Vincent Corvinelli, and Calisto Zuzarte. 2015. Cardinality Estimation using Neural Networks. In CASCON. 53–59.
- [40] Hoang Vu Nguyen, Klemens Böhm, Florian Becker, Bertrand Goldman, Georg Hinkel, and Emmanuel Müller. 2015. Identifying User Interests within the Data Space-a Case Study with SkyServer.. In EDBT. 641–652.
- [41] William O'Mullane, Nolan Li, María Nieto-Santisteban, Alex Szalay, Ani Thakar, and Jim Gray. 2005. Batch is Back: CasJobs, Serving Multi-TB Data on the Web. In *ICWS*. 33–40.
- [42] Terence Parr. 2013. The Definitive ANTLR 4 Reference. Pragmatic Bookshelf.
- [43] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-Learn: Machine Learning in Python. *JMLR* 12 (2011), 2825–2830.
- [44] M Jordan Raddick, Ani R Thakar, Alexander S Szalay, and Rafael DC Santos. 2014. Ten Years of SkyServer I: Tracking Web and SQL e-Science Usage. Computing in Science & Engineering 16, 4 (2014), 22–31.
- [45] M Jordan Raddick, Ani R Thakar, Alexander S Szalay, and Rafael DC Santos. 2014. Ten Years of SkyServer II: How Astronomers and the Public Have Embraced e-Science. *Computing in Science & Engineering* 16, 4 (2014), 32–40.
- [46] Yanyan Shen, Kaushik Chakrabarti, Surajit Chaudhuri, Bolin Ding, and Lev Novik. 2014. Discovering Queries Based on Example Tuples.

In SIGMOD. 493-504.

- [47] Vik Singh, Jim Gray, Ani Thakar, Alexander S Szalay, Jordan Raddick, Bill Boroski, Svetlana Lebedeva, and Brian Yanny. 2007. Skyserver Traffic Report - The First Five Years. ArXiv 0701173 (2007).
- [48] Kostas Stefanidis, Marina Drosou, and Evaggelia Pitoura. 2009. "You May Also Like" Results in Relational Databases. In *PersDB*. 37–42.
- [49] Alexander S Szalay. 2018. From SkyServer to SciServer. AAPSS 675, 1 (2018), 202–220.
- [50] Alexander S Szalay, Jim Gray, Ani R Thakar, Peter Z Kunszt, Tanu Malik, Jordan Raddick, Christopher Stoughton, and Jan vandenBerg. 2002. The SDSS Skyserver: Public Access to the Sloan Digital Sky Server Data. In *SIGMOD*. 570–581.
- [51] Kai Sheng Tai, Richard Socher, and Christopher D Manning. 2015. Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks. ArXiv 1503.00075 (2015).
- [52] TPC. 2018. TPC Benchmarks. http://www.tpc.org.
- [53] Wei Wang, Meihui Zhang, Gang Chen, H. V. Jagadish, Beng Chin Ooi, and Kian-Lee Tan. 2016. Database Meets Deep Learning: Challenges and Opportunities. *SIGMOD Rec.* 45, 2 (2016), 17–22.
- [54] Wenpeng Yin, Katharina Kann, Mo Yu, and Hinrich Schütze. 2017. Comparative Study of CNN and RNN for Natural Language Processing.

ArXiv 1702.01923 (2017).

- [55] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. 2014. How Transferable Are Features in Deep Neural Networks?. In NIPS. 3320– 3328.
- [56] Tom Young, Devamanyu Hazarika, Soujanya Poria, and Erik Cambria. 2018. Recent Trends in Deep Learning Based Natural Language Processing. *IEEE Computational Intelligence Magazine* 13, 3 (2018), 55–75.
- [57] Wojciech Zaremba and Ilya Sutskever. 2014. Learning to Execute. ArXiv 1410.4615 (2014).
- [58] Zhong Zeng, Zhifeng Bao, Thuy Ngoc Le, Mong Li Lee, and Wang Tok Ling. 2014. ExpressQ: Identifying Keyword Context and Search Target in Relational Keyword Queries. In *CIKM*. 31–40.
- [59] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-Level Convolutional Networks for Text Classification. In NIPS. 649–657.
- [60] Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2SQL: Generating Structured Queries From Natural Language using Reinforcement Learning. ArXiv 1709.00103 (2017).
- [61] Zainab Zolaktaf, Mostafa Milani, and Rachel Pottinger. 2020. Facilitating SQL Query Composition and Analysis. ArXiv 2002.09091 (2020).