

Database-Agnostic Workload Management

Shrainik Jain

University of Washington
shrainik@cs.washington.edu

Thierry Cruanes

Snowflake Computing
thierry.cruanes@snowflake.net

Jiaqi Yan

Snowflake Computing
jiaqi.yan@snowflake.net

Bill Howe

University of Washington
billhowe@cs.washington.edu

ABSTRACT

We present a system to support generalized SQL workload analysis and management for multi-tenant and multi-database platforms. Workload analysis applications are becoming more sophisticated to support database administration, model user behavior, audit security, and route queries, but the methods rely on specialized feature engineering, and therefore must be carefully implemented and reimplemented for each SQL dialect, database system, and application. Meanwhile, the size and complexity of workloads are increasing as systems centralize in the cloud. We model workload analysis and management tasks as variations on query labeling, and propose a system design that can support general query labeling routines across multiple applications and database backends. The design relies on the use of learned vector embeddings for SQL queries as a replacement for application-specific syntactic features, reducing custom code and allowing the use of off-the-shelf machine learning algorithms for labeling. The key hypothesis, for which we provide evidence in this paper, is that these learned features can outperform conventional feature engineering on representative machine learning tasks. We present the design of a database-agnostic workload management and analytics service, describe potential applications, and show that separating workload representation from labeling tasks affords new capabilities and can outperform existing solutions for representative tasks, including workload sampling for index recommendation and user labeling for security audits.

ACM Reference Format:

Shrainik Jain, Jiaqi Yan, Thierry Cruanes, and Bill Howe. 2019. Database-Agnostic Workload Management. In *Proceedings of Conference on Innovative Data Systems Research (CIDR'19)*. ACM, New York, NY, USA, Article 4, 6 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Extracting patterns from a SQL query workload has enabled a number of important features in database systems, including workload compression [3], index recommendation [2], modeling user and application behavior [9, 31, 35], query recommendation [1], predicting cache performance [5, 29], and designing benchmarks [35].

These techniques can be used as part of a more comprehensive approach to automate database administration [26].

However, the diversity of applications have led to a diversity of solutions, each relying on specialized feature engineering. For example, workload summarization for index recommendation uses the structure of join and group by operators as features [3], query recommendation may pre-process a query into fragments before making recommendations [13], and security audits may require user-defined functions to enforce particular policies [32].

In fact, the features and the algorithms to extract them tend to be the significant contributions in the papers in this space. But the state of the art in a variety of applications is to learn features automatically. For instance, Natural Language Processing applications previously relied on parsing and labeling sentences as a pre-processing step, but now use learned vector representations almost exclusively [6, 28]. This approach not only obviates the need for manual feature engineering and pre-processing, but also has the potential to significantly outperform more specialized methods.

We see three trends motivating an analogous role for generalized workload representations. First, workload heterogeneity is increasing, making it difficult to maintain SQL parsers and feature extraction routines. The number of SQL-like languages is increasing, with inconsistent support and syntax for even relatively common features such as outer joins. Second, workload scale is increasing. Cloud-hosted, multi-tenant database services including Redshift [8], Snowflake [4], BigQuery [23] and more receive millions of queries daily from thousands of customers using hundreds of schemas; relying on brittle parsers (or worse, manual inspection) to identify query patterns that influence administration decisions is no longer tenable. Third, new use cases for centralized workload management are emerging. For example, SQL debugging [7], database forensics [27], and data use management [32] motivate a more automated analysis of user behavior patterns, and cloud-hosted multi-tenant systems motivate a more automated approach to query routing and resource allocation.

In this work, we propose Querc, a database-agnostic systems for mining and managing large-scale and heterogeneous workloads. We model workload management and analysis as a set of query labeling tasks. For instance, workload sampling can be reduced to labeling each query as present or absent in the sample, error prediction involves labeling each query with an error type, query routing involves labeling each query with a cluster resource to which the query should be routed, and so on. Because our framework depends only on the query text (along with typical metadata such as arrival timestamp and userid issuing the query), it can be used with any DBMS and any SQL dialect. In fact, as we will show, features learned

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CIDR'19, Jan 2019, Asilomar, California USA

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

with a workload against a particular schema and SQL dialect can be effective even when used with a *different* schema and SQL dialect.

The weakness of this approach is that it requires enormous amounts of data to be effective. But as database products migrate to the cloud, service providers have access to workloads from a large number of customers, potentially even across different database products. Since the input is just the query text, these diverse workloads can be processed as one very large dataset. But the resulting vectors can still be used to train models to support specific applications, as we will show on two representative tasks: workload summarization for index selection and user prediction for security audits and routing.

2 SYSTEM ARCHITECTURE

Figure 1 illustrates the architecture of Querc. There are three applications, X, Y, Z. Each application has its own database, DB(X), DB(Y), and DB(Z), though these may be logical instances in the same physical multi-tenant service. In this example, DB(X) and DB(Y) are tenants in the same service. Each application is also associated with a separate stream of queries (at left), where query(X,t) indicates a batch of queries arriving for application X at time instant t.

Each application is associated with one Qworker, but each Qworker operates multiple classifiers. Qworkers may not be entirely stateless, as some labeling tasks process a small window of queries. However, the state is assumed to be small such that the Qworkers do not need their own local storage and can be load balanced and parallelized in typical ways. Each classifier is a pre-trained (embedder, labeler) pair. The same trained embedder may be used across multiple applications. This split design is critical, because we want to learn features using a very large, combined workload, but an individual classifier may perform better when trained on an application-specific workload. In this example, application X and application Y both share the same embedder, EmbedderA, trained on the combined X and Y workloads, written EmbedderA(X,Y). This log sharing between customers may not always be permitted by customers for security reasons, and in this example, application Z uses only its own data. But there is some incentive for customers to pool their data as the additional signal can potentially improve accuracy, and some cloud providers support features to allow data sharing between customers.

The Labeler passes the query on to the database, but also transmits the query back to a central training module (“Training, Evaluation, and Offline Labeling” in Figure 1). The training module manages training sets, including the (parallel) execution of training and evaluation routines, then deploys trained models back to Qworkers. There is significant ongoing research in the database, systems, and ML communities on runtime architectures for training and deploying models (e.g., [21]); we do not discuss them further since our requirements are relatively modest.

Since Querc is specialized for query workload analytics rather than general machine learning, one data model can be shared among most applications. The only messages passed between components are labeled queries. A labeled query is a tuple $(Q, c_1, c_2, c_3, \dots)$ where c_i is a label. This simple model captures situations where a query arrives already equipped with a timestamp, a userid, an IP

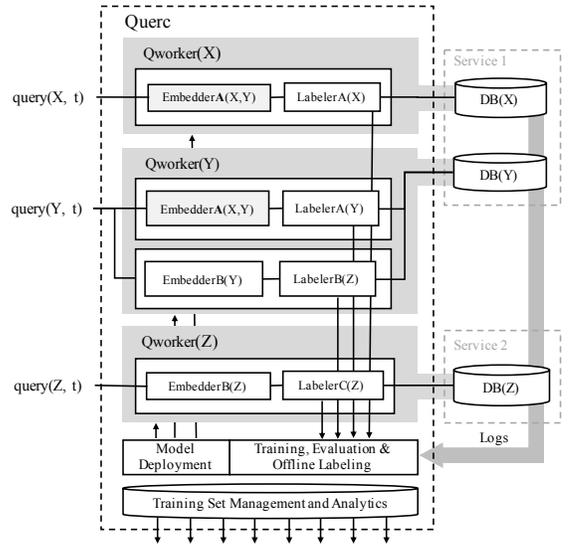


Figure 1: System architecture. Queries arrive for three different applications X, Y, and Z and are processed by one or more (embedder, labeler) pair before being sent on to the database, centralized for offline labeling tasks, or both.

address, etc., but also captures more verbose query logs that are returned from the database.

The training module also records the queries with their predicted labels for retraining, evaluation, and to support offline analysis tasks. Offline tasks are those that do not require or do not allow processing each query separately, and can be implemented as typical batch jobs. For example, query clustering is important for workload summarization [16], but does not require real-time labeling of individual queries.

Training data is collected periodically from the databases in the form of query logs. These logs are (batched) sequences of labeled queries, but with additional labels to be used for training, such as runtime, memory usage, error codes, security flags, resource IDs. We do not specify the mechanism by which these logs are transmitted from the database to Querc, since most systems have robust means of exporting logs in appropriate forms.

In some applications, Querc may not be in the critical path for query execution to avoid any performance overhead or reduce dependencies. In these cases, queries will be forked to Querc. No change to the architecture is required in this case; queries come in, and labeled queries are collected in the training module. The query is simply not forwarded to the database.

This architecture is not designed for continuous learning, as the training is handled separately from real time query labeling. Not all algorithms can support fully continuous learning, and an important design goal is to support simple machine learning algorithms as labelers. Model training is therefore assumed to occur infrequently as a batch job.

3 LEARNING VECTOR REPRESENTATIONS

There are multiple choices for embedders; we describe two initial models we evaluate in this paper:

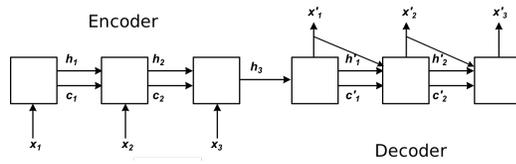


Figure 2: The LSTM Autoencoder network architecture learns to generate the input token in the decoding phase. Once trained, the encoder can be used to output a vector representation for the text of a query.

Context prediction models: Mikolov et al. [17, 24, 25] proposed learning a vector representation for words by predicting the next word in a context, and then deriving a vector representation for larger semantic units (sentences, paragraphs, documents) by adding a vector representing the paragraph to each context as an additional “word.” The learned vector for this virtual context word is used as a representation for the entire paragraph. This “Doc2Vec” method has been shown to capture semantic relationships that work well for, say, sentiment classification and clustering tasks [14, 18]. This approach can be applied directly for learning representations of SQL queries: We can use fixed-size context windows to learn a representation for each token in the query, and include an identifier to learn a representation of entire query.

LSTM AutoEncoders: The paragraph vector approach in the previous section is viable, but it requires a hyper-parameter for the context size. There is no obvious way to determine a context size for queries, for two reasons: First, there may be semantic relationships between distant tokens in the query. Second, the length of queries vary widely in ad hoc workloads [10, 12]. To avoid setting a context size, we can use Long Short-Term Memory (LSTM) networks [36], which are modified Recurrent Neural Networks (RNN) that can automatically learn how much context to remember and how much of it to forget, thereby removing the dependency on a fixed context size. LSTMs have successfully been used in sentence classification, semantic similarity between sentences and sentiment analysis [30]. We use a standard LSTM encoder decoder network [20, 37] with architecture as illustrated in Figure 2.

An LSTM autoencoder is trained by sequentially feeding words from the query to the network one word at a time, and then attempting to reproduce the input. The LSTM network not only learns the encoding for the samples, but also the relevant context window associated with the samples. The final output of the encoder network gives us an encoding for the query. Once this network has been trained, an embedded representation for a query can be computed by passing the query to the encoder network, completing a forward pass, and using the hidden state of the final encoder LSTM cell as the learned vector representation.

There are multiple prior approaches in the NLP literature that compare the efficacy of these models and their relative performance [17, 22, 30]. For this paper, we consider context-based models (i.e., doc2vec) and LSTM AutoEncoders.

4 APPLICATIONS

The applications supported by this system reduce to *query labeling*, and general workflow consists of two machine learning models: a

representation learner (an embedder) and a classifier. We split the task into two parts to allow the same representation to be used for multiple applications.

Workload summarization for index recommendation: The goal [3, 16] is to find a representative sample of the workload as input to further database administration, tuning, and testing tasks [3, 33]. In particular, workload summarization aids index recommendation, since the recommendation process is typically quadratic in the size of the workload [3]. While index recommendation systems are well-studied and ship with most production databases [2, 3], the quality of the representative sample determines the overall quality of the final recommendations. In Section 5, we show that a simple sampling procedure using learned features delivers a significant runtime improvement over the built-in sampling procedure in the SQL Server database system.

Enforcing query routing policies: Query Routing in a distributed database involves identifying the cluster resources on which to execute the incoming query. The policies that govern these routing decisions may involve customer SLAs, security considerations (e.g., certain applications must use a physically distinct cluster from other applications), auditing requirements (e.g., queries from certain accounts or those accessing certain tables must be logged for auditing purposes). Even in modern cloud-hosted database products such as Snowflake [4] and BigQuery [23], these policies tend to be manually encoded, and management of these policies as they evolve, while maintaining multiple heterogeneous clusters used by thousands of customers, is increasingly perceived as untenable. Under the hypothesis that queries that follow a particular policy tend to have similar features, Querc can help identify policy misconfiguration by detecting when a predicted routing decision differs from the assigned routing decision.

Error prediction: Particular syntax patterns in the workload may be associated with resource errors or bugs in the database system. In a multi-tenant, multi-database, and high-volume scenario, identification of the syntactic patterns that tend to trigger errors, either manually or with scripts, becomes untenable: there may be hundreds of error codes, each with hundreds of subtle patterns that tend to trigger them, across hundreds of tenant schemas. Using learned features, a classifier to predict errors from syntax is trivial to engineer. This prediction allows the query to be routed to a different runtime environment that is instrumented, equipped with more memory per node, or running a more stable version of the database engine. We consider this application in a tech report companion to this paper [11].

Resource allocation: The structure of the query is not sufficient to accurately predict its runtime or memory footprint, but it can provide a hint that can be used for load balancing, scheduling, and as an input for optimization. If we can coarsely categorize queries as memory-intensive, long-running, etc. with some degree of accuracy, these labels can be used as a simple, database-agnostic way to speculatively allocate resources. Training data is readily available from the query logs themselves. We consider this application in a tech report companion to this paper [11].

Query recommendation: The query recommendation problem can be modeled as a prediction of the next query the user will submit to the database based on the recent history of queries [1]. This prediction is then shown to the user through an appropriate

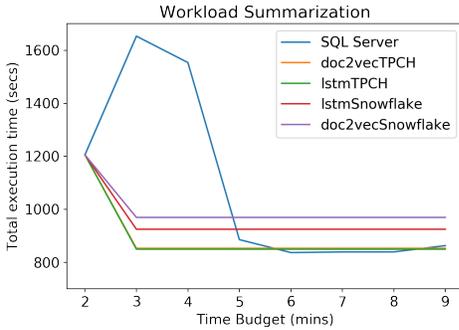


Figure 3: Workload runtime using indexes recommended under various time budgets. For most time budgets, the workload summaries improve runtimes, even when the embedders were trained on an unrelated workload (lstmSnowflake and doc2vecSnowflake).

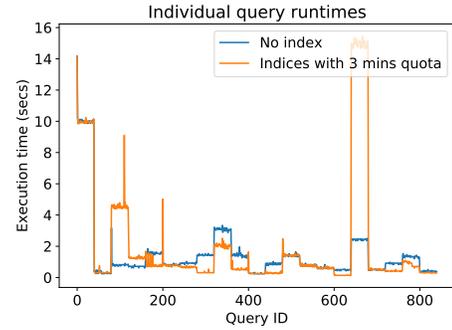


Figure 4: Runtime for each query under no indexes and under indexes recommended with a three-minute time budget. For a few specific queries, the presence of a recommended index results in significantly worse performance.

client application to assist in query authoring. Our framework can generate features that can be used to train query recommendation models. We consider this application in a tech report companion to this paper [11].

Security auditing: To the extent that users’ individual workloads tend to follow predictable patterns, an anomalous query may be a sign that a user’s account has been compromised. By formulating a prediction problem that tries to guess the user that submitted the query from the syntax alone, we can identify anomalous queries for security audits. In our framework, the labeler is a simple classifier $V \rightarrow user$.

5 EXPERIMENTS

We consider two applications: Workload summarization for index selection, and labeling tasks for security audits and query routing.

5.1 Workload Summaries for Index Selection

The workload summarization task (with respect to index recommendation) is to find a subset Q_{sub} of a given query workload Q , such that the set of indexes recommended based on Q_{sub} is similar to the set of indexes recommended for the overall workload Q . Previous solutions are primarily variants of the approach of Chaudhuri et al. [3], which uses K-medoids to cluster the queries and selects a witness query from each cluster. However, the authors emphasize that a custom distance function should be developed for specific workloads; our hypothesis is that generic representation learning approaches obviate the need for these custom distance functions.

In the Querc framework, this task is offline and does not require real-time labeling of queries. Instead, we perform the task as an offline unsupervised learning task. In our approach, we assign each query to a vector (using a suitably trained embedder), then simply use K-means to find K query clusters and pick the nearest query to the centroid in each cluster as the representative subset. To determine K , we use an intentionally simple method (the “elbow method” [15]) which runs the K-means algorithm in a loop with increasing

K till the rate of change of the sum of squared distances from centroids plateaus. Although better methods exist, we highlight the effect of the learned vectors rather than the choice of K .

Setup: Following the evaluation strategy of Chaudhuri et al. [3], we first run the index selection tool on the entire workload Q , create the recommended indexes, and measure the runtime t_{orig} for the original workload. We then run use the workload summarization algorithm to produce a reduced set of queries Q_{sub} , re-run the index selection tool, create the recommended indexes, and again measure the runtime t_{sub} of the entire original workload. We use SQL Server 2016 and the Database Engine Tuning Advisor, which performs its own summarization on the input according to the documentation. We use an *m4.large* AWS EC2 instance as the server. We use TPC-H with scale factor 1 as the workload for comparison with previous results and to interpret the recommended indexes, but we also show how the method performs when trained on a more complex Snowflake workload.

We pass the summarized workload to the tuning advisor, along with a time budget (a parameter supported by the tuning advisor). Each experiment involves clearing caches, generating indexes, applying the indexes, and running the full workload. We report the time running the workload; the time budget specifies the time limit under which the advisor must return a set of recommendations.

Results: Figure 3 shows the results. The x-axis is the time budget, and the y-axis is the runtime for the entire workload after building the recommended indexes. For time budgets less than 3 minutes, the advisor does not produce any index recommendations for any method, and the runtime is constant at 1200 seconds. As we relax the time budget, different sets of indexes are recommended, each associated with a separate runtime. The full workload (blue line) varies dramatically with the time budget, and surprisingly it gets worse before it gets better. For the summarized workloads, the workload is small enough that the runtimes are constant: Once three minutes have elapsed, the advisor has found the “optimal” set of indexes, and allowing more time does not change the result.

We evaluate four trained embedders: two methods on two workloads. The two methods are Doc2Vec and the LSTMAutoencoder, and the two workloads are TPC-H itself, and a separate workload

	Account Labeling	User Labeling
Doc2Vec	78.8%	39%
LSTMAutoencoder	99.1%	55.4%

Table 1: Query Labeling results

of 500,000 queries from the Snowflake service. When training the embedder on TPC-H (doc2VecTPCH and lstmTPCH), the advisor finds close-to-optimal indexes in about three minutes as opposed to the six minutes the advisor requires on the full workload. Surprisingly, under tight time budgets, the index recommendations made by the native system can actually *hurt* performance relative to having no indexes at all. The reason is that the optimizer chooses a bad plan for a few particular queries, but the effect is enough to hurt the overall runtime. Figure 4 shows the sequence of queries in the workload on the x-axis and the runtime for each query on the y-axis under no indexes and the low-quality indexes found at the three-minute time budget. All instances of TPC-H query 18 (queries 640-680 in Figure 4) take much longer than they would take when run without these indexes, because the optimizer finds a bad plan.

Transfer Learning: Figure 3 also illustrates the capacity for transfer learning using Querc: When training the embedder on the snowflake dataset – a completely unrelated workload to TPC-H workload in the *SQL Server dialect* – the summarized workload still outperforms native SQL Server for most time budgets. This transfer learning effect allows us to bootstrap new applications without waiting for a representative workload to accumulate, and to avoid having to repeatedly re-implement brittle parsers and feature extractors for each new dialect of SQL we encounter.

5.2 Labeling for Security Audits

We consider the conditions under which the learned features from query syntax are sufficient to predict username and customer account, where each customer has many users. When the predicted username differs from the actual username, we can potentially flag the query for an audit. Predicting username can help flag queries for security audits, account and cluster labels can identify misrouted queries. Labels from query syntax using the two embedding methods described in Section 3 over the Snowflake dataset.

Setup: We use embedders pre-trained on 500000 Snowflake queries. The experiment itself is run on another dataset of 200000 Snowflake queries labeled with username, account_id and cluster_name for the cluster that ran the query. Next we train classifiers (randomized decision trees) for username and customer account.

Results: Table 1 shows the results for the labeling experiments. The numbers denote the 10-fold cross validation score on the respective task. We find that LSTM based embedders beats Doc2Vec on all tasks. The LSTM method achieves near perfect accuracy when predicting the customer account, which is because it automatically incorporates signal from the schema, and different customers use primarily different schemas (there are instances of shared schemas, but that is the less common case). The method was completely generic and knows nothing about schemas or queries. For user prediction, the task is more difficult, and the overall accuracy is lower at 55%. Upon further analysis we found that the user labeling

#queries	#users	accuracy
73881	28	49.3%
55333	10	37.4%
18487	46	31.8%
5471	21	96.2%
4213	6	58.5%
3894	12	99.7%
3373	9	99.8%
2867	6	99.8%
1953	15	89.1%
1924	4	98.1%
1776	9	95.2%
1699	5	99.8%
1108	12	98.2%

Table 2: Top accounts with user prediction accuracy.

task has > 95% accuracies for a majority of accounts (Table 2). The accounts that had poor accuracies for user labeling had one distinctive property: multiple users running the exact same query, making the users nearly indistinguishable. In the sample of workload that we were working with, there were two accounts that had a number of repetitive queries by different users (for instance, 69% percent of the 74000 queries in an account had more than one user label), and these two accounts also covered around 65% of the total queries, bringing down the overall accuracy of classifiers.

6 FUTURE WORK

Other methods: There are a variety of other methods for learning representations of text that we do not evaluate in this paper. Our goal is not to identify the best possible representation learning approach but rather to show that these methods can compete with and outperform classical approaches that rely on task-specific heuristics and feature engineering (extracting JOIN clauses, counting the number of attributes, etc.), and to organize the methods into a coherent system architecture.

Alternative methods can be roughly categorized into *non-neural-network* based methods and *neural-network*-based methods. The non-neural-network-based methods, including non-negative matrix factorization (NMF), bag-of-words representations, and LDA [22] have been shown to be less effective than neural-network-based-methods in a variety of contexts [19, 25]. Apart from the methods considered in this paper, there are more recent neural-network-based methods using Convolutional Neural Networks (CNNs) adapted for text data. However, Yin et al. [34] showed that RNN based methods (e.g., LSTMs) perform well and are robust in a broad range of tasks when compared to CNNs. However, we plan to extend the current work to include a rigorous comparison of the techniques not covered in this paper.

Publish pre-trained models: The results in Section 5 demonstrate that the proposed framework in this paper has potential to use pre-trained models on generic workloads to aid analytics for previously unseen query. In future work, we will build this framework as a service which is accessible by third parties. Given the workloads that we have access to from Snowflake [4], such a service could be

really beneficial for researchers who do not have access to massive query workloads.

7 CONCLUSIONS

We presented the architecture for Querc, a database-agnostic workload analytics service that captures the structural and schema patterns present in the query workload automatically, largely eliminating the need for the specialized syntactic feature engineering that has motivated a number of papers in the literature. The proposed architecture provides a new way of organizing a variety of database administration and user productivity tasks, and provides a mechanism by which to automatically adapt database operations to specific query workloads. Our evaluation of this architecture showed that our general framework outperformed or was competitive with previous approaches that required specialized feature engineering, and also admitted simpler classification algorithms because the inputs are numeric vectors with well-behaved algebraic properties rather than result of arbitrary user-defined functions for which few properties can be assumed. The use of transfer learning in Querc allows workload analytics to be *SQL dialect independent* and enables the capability to bootstrap new analytics tasks and avoid re-implementing brittle codes paths.

REFERENCES

- [1] Javad Akbarnejad, Gloria Chatzopoulou, Magdalini Eirinaki, Suju Koshy, Sarika Mittal, Duc On, Neoklis Polyzotis, and Jothi S Vindhiya Varman. 2010. SQL QueRIE recommendations. *Proceedings of the VLDB Endowment* 3, 1-2 (2010), 1597–1600.
- [2] Surajit Chaudhuri, Prasanna Ganesan, and Vivek Narasayya. 2003. Primitives for Workload Summarization and Implications for SQL. In *Proceedings of the 29th International Conference on Very Large Data Bases - Volume 29 (VLDB '03)*. VLDB Endowment, 730–741. <http://dl.acm.org/citation.cfm?id=1315451.1315514>
- [3] Surajit Chaudhuri, Ashish Kumar Gupta, and Vivek Narasayya. 2002. Compressing sql workloads. In *Proceedings of the 2002 ACM SIGMOD international conference on Management of data*. ACM, 488–499.
- [4] Benoit Dageville, Thierry Cruanes, Marcin Zukowski, Vadim Antonov, Artin Avanes, Jon Bock, Jonathan Claybaugh, Daniel Engovatov, Martin Hentschel, Jiansheng Huang, Allison W. Lee, Ashish Motivala, Abdul Q. Muir, Steven Pelley, Peter Povinec, Greg Rahn, Spyridon Triantafyllis, and Philipp Unterbrunner. 2016. The Snowflake Elastic Data Warehouse. In *Proceedings of the 2016 International Conference on Management of Data (SIGMOD '16)*. ACM, New York, NY, USA, 215–226. <https://doi.org/10.1145/2882903.2903741>
- [5] Asit Dan, Philip S. Yu, and Jen Yao Chung. 1995. Characterization of Database Access Pattern for Analytic Prediction of Buffer Hit Probability. *The VLDB Journal* 4, 1 (Jan. 1995), 127–154. <https://doi.org/10.1007/BF01232474>
- [6] Yoav Goldberg and Omer Levy. 2014. word2vec explained: Deriving mikolov et al.'s negative-sampling word-embedding method. *arXiv preprint arXiv:1402.3722* (2014).
- [7] Torsten Grust and Jan Rittinger. 2013. Observing SQL Queries in Their Natural Habitat. *ACM Trans. Database Syst.* 38, 1, Article 3 (April 2013), 33 pages. <https://doi.org/10.1145/2445583.2445586>
- [8] Anurag Gupta, Deepak Agarwal, Derek Tan, Jakub Kulesza, Rahul Pathak, Stefano Stefani, and Vidhya Srinivasan. 2015. Amazon redshift and the case for simpler data warehouses. In *Proceedings of the 2015 ACM SIGMOD international conference on management of data*. ACM, 1917–1923.
- [9] Shrainik Jain and Bill Howe. 2016. Data Cleaning in the Wild: Reusable Curation Idioms from a Multi-Year SQL Workload. In *Proceedings of the 11th International Workshop on Quality in Databases, QDB 2016, at the VLDB 2016 conference, New Delhi, India, September 5, 2016*.
- [10] Shrainik Jain and Bill Howe. 2016. SQLShare Data Release. https://uwescience.github.io/sqlshare/data_release.html. [Online;].
- [11] Shrainik Jain and Bill Howe. 2018. Query2Vec: NLP Meets Databases for Generalized Workload Analytics. *CoRR* abs/1801.05613 (2018). [arXiv:1801.05613](https://arxiv.org/abs/1801.05613)
- [12] Shrainik Jain, Dominik Moritz, Daniel Halperin, Bill Howe, and Ed Lazowska. 2016. SQLShare: Results from a Multi-Year SQL-as-a-Service Experiment. In *Proceedings of the 2016 International Conference on Management of Data (SIGMOD '16)*. ACM, 281–293. <https://doi.org/10.1145/2882903.2882957>
- [13] Nodira Khoussainova, YongChul Kwon, Magdalena Balazinska, and Dan Suciu. 2010. SnipSuggest: Context-Aware Autocompletion for SQL. *PVLDB* 4, 1 (2010), 22–33. <http://www.vldb.org/pvldb/vol4/p22-khoussainova.pdf>
- [14] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. *CoRR* abs/1408.5882 (2014). <http://arxiv.org/abs/1408.5882>
- [15] Trupti M Kodinariya and Prashant R Makwana. 2013. Review on determining number of Cluster in K-Means Clustering. *International Journal* 1, 6 (2013), 90–95.
- [16] Piotr Kolaczowski. 2008. Compressing very large database workloads for continuous online index selection. In *Database and Expert Systems Applications*. Springer, 791–799.
- [17] Quoc V Le and Tomas Mikolov. [n. d.]. Distributed Representations of Sentences and Documents.
- [18] Yann LeCun. [n. d.]. THE MNIST DATABASE of handwritten digits. <http://yann.lecun.com/exdb/mnist/>
- [19] Omer Levy, Yoav Goldberg, and Israel Ramat-Gan. 2014. Linguistic Regularities in Sparse and Explicit Word Representations. In *CoNLL*. 171–180.
- [20] Jiwei Li, Minh-Thang Luong, and Dan Jurafsky. 2015. A Hierarchical Neural Autoencoder for Paragraphs and Documents. *CoRR* abs/1506.01057 (2015). [arXiv:1506.01057](https://arxiv.org/abs/1506.01057)
- [21] Mu Li, David G. Andersen, Jun Woo Park, Alexander J. Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J. Shekita, and Bor-Yiing Su. 2014. Scaling Distributed Machine Learning with the Parameter Server. In *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation (OSDI'14)*. USENIX Association, Berkeley, CA, USA, 583–598. <http://dl.acm.org/citation.cfm?id=2685048.2685095>
- [22] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies-volume 1*. Association for Computational Linguistics, 142–150.
- [23] Sergey Melnik, Andrey Gubarev, Jing Jing Long, Geoffrey Romer, Shiva Shivakumar, Matt Tolton, and Theo Vassilakis. 2010. Dremel: interactive analysis of web-scale datasets. *Proceedings of the VLDB Endowment* 3, 1-2 (2010), 330–339.
- [24] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [25] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [26] Andrew Pavlo, Gustavo Angulo, Joy Arulraj, Haibin Lin, Jiexi Lin, Lin Ma, Prashanth Menon, Todd C Mowry, Matthew Perron, Ian Quah, et al. [n. d.]. Self-Driving Database Management Systems.
- [27] Kyriacos E. Pavlou and Richard T. Snodgrass. 2013. Generalizing Database Forensics. *ACM Trans. Database Syst.* 38, 2, Article 12 (July 2013), 43 pages. <https://doi.org/10.1145/2487259.2487264>
- [28] Jeffrey Pennington, Richard Socher, and Christopher D Manning. [n. d.]. Glove: Global Vectors for Word Representation.
- [29] Carsten Sapia. 2000. PROMISE: Predicting Query Behavior to Enable Predictive Caching Strategies for OLAP Systems. In *Proceedings of the Second International Conference on Data Warehousing and Knowledge Discovery (DaWaK 2000)*. Springer-Verlag, London, UK, UK, 224–233. <http://dl.acm.org/citation.cfm?id=646109.679288>
- [30] Duyu Tang, Bing Qin, and Ting Liu. [n. d.]. Document Modeling with Gated Recurrent Neural Network for Sentiment Classification.
- [31] Quoc Trung Tran, Konstantinos Morfonios, and Neoklis Polyzotis. 2015. Oracle Workload Intelligence. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD '15)*. ACM, New York, NY, USA, 1669–1681. <https://doi.org/10.1145/2723372.2742791>
- [32] Prasang Upadhyaya, Magdalena Balazinska, and Dan Suciu. 2015. Automatic Enforcement of Data Use Policies with DataLawyer. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD '15)*. ACM, New York, NY, USA, 213–225. <https://doi.org/10.1145/2723372.2723721>
- [33] Jiaqi Yan, Qiuye Jin, Shrainik Jain, Stratis D. Viglas, and Allison Lee. 2018. Snow-trail: Testing with Production Queries on a Cloud Database. In *Proceedings of the Workshop on Testing Database Systems (DBTest '18)*. ACM, New York, NY, USA, Article 4, 6 pages. <https://doi.org/10.1145/3209950.3209958>
- [34] Wenpeng Yin, Katharina Kann, Mo Yu, and Hinrich Schütze. 2017. Comparative study of cnn and rnn for natural language processing. *arXiv preprint arXiv:1702.01923* (2017).
- [35] Philip S. Yu, Ming-Syan Chen, Hans-Ulrich Heiss, and Sukho Lee. 1992. On Workload Characterization of Relational Database Environments. *IEEE Trans. Softw. Eng.* 18, 4 (April 1992), 347–355. <https://doi.org/10.1109/32.129222>
- [36] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329* (2014).
- [37] Richard S. Zemel. 1994. Autoencoders, minimum description length and Helmholtz free energy. NIPS.