Efficient Deep Learning Pipelines for Accurate Cost Estimations Over Large Scale Query Workload

Johan Kok Zhi Kang
johan.kok@u.nus.edu
National University of Singapore
Singapore

Gaurav
gaurav@grab.com
GrabTaxi Holdings
Singapore

Sien Yi Tan
sienyi.tan@grab.com
GrabTaxi Holdings
Singapore

Feng Cheng
feng.cheng@grab.com
GrabTaxi Holdings
Singapore

Shixuan Sun
sunsx@comp.nus.edu.sg
National University of Singapore
Singapore

Bingsheng He
hebs@comp.nus.edu.sg
National University of Singapore
Singapore

ABSTRACT

The use of deep learning models for forecasting the resource consumption patterns of SQL queries have recently been a popular area of study. While these models have demonstrated promising accuracy, training them over large scale industry workloads are expensive. Space inefficiencies of encoding techniques over large numbers of queries and excessive padding used to enforce shape consistency across diverse query plans implies 1) longer model training time and 2) the need for expensive, scaled up infrastructure to support batched training. In turn, we developed *Prestroid*, a tree convolution based data science pipeline that accurately predicts resource consumption patterns of query traces, but at a much lower cost. We evaluated our pipeline over 19K Presto OLAP queries, on a data lake of more than 20PB of data from Grab. Experimental results imply that our pipeline outperforms benchmarks on predictive accuracy, contributing to more precise resource prediction for large-scale workloads, yet also reduces per-batch memory footprint by 13.5x and per-epoch training time by 3.45x. We demonstrate direct cost savings of up to 13.2x for large batched model training over Microsoft Azure VMs.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**; • **Information systems** → *Information retrieval query processing*.

KEYWORDS

Tree Convolution; Query resource allocation; Deep Learning;

ACM Reference Format:

Johan Kok Zhi Kang, Gaurav, Sien Yi Tan, Feng Cheng, Shixuan Sun, and Bingsheng He. 2021. Efficient Deep Learning Pipelines for Accurate Cost Estimations Over Large Scale Query Workload. In *Proceedings of the 2021 International Conference on Management of Data (SIGMOD '21)*, June 20–25, 2021, Virtual Event, China. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3448016.3457546>

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).

SIGMOD '21, June 20–25, 2021, Virtual Event, China

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8343-1/21/06.

<https://doi.org/10.1145/3448016.3457546>

1 INTRODUCTION

Figure 1: Integration of Prestroid for resource deployment

Present trends in big (OLAP) data query engine design have shown two key features; compatibility with cloud-based infrastructure and the adoption of a decoupled compute-storage paradigm. This shift introduces complexities that makes the use of analytical models to forecast a query's expected resource consumption increasingly difficult. In turn, there has been a rising adoption in using deep learning models [17, 21, 26] for such a task. These models, when trained over features extracted from prior executed queries, can yield good results in predicting the resource consumption of new, unseen queries.

Our work is heavily inspired by use cases from Grab, a large superapp company managing a data lake of more than 20PB in the cloud. As the business grows and diversifies, queries to our data lake are expanding both in volume and in variety. We record hundreds of thousands of queries issued over multiple Presto [33] clusters powered by cloud infrastructure. Monthly cloud expenditures form a large portion of expenses, thus having a systematic and accurate framework for forecasting a query's resource utilization will entail huge cost savings [10, 11]. Such a framework may be represented as an end-to-end pipeline in Fig 1. Incoming queries are first parsed for their features before being routed through a deep learning model. The model predicts the resources needed by the query and these resources are created for query execution. Such a system assures that resources allocated to run a query are neither excessive which incurs cost, nor insufficient at the risk of a query violating its service level agreement (SLA).

To date, researchers have explored various deep learning pipelines for query-resource prediction. We review those works in Section 2.

Herein, we observed two issues with a direct application of these pipelines. The first is that pipeline models are costly to be trained over the query patterns from Grab, in large variety and volume. Our experiments showed that a single model training over 19K queries can cost as much as \$76.25 for batch sizes of 256. The second is that model updates need to be done in perpetuity to keep up with the rapid evolution of the business. This amounts to a spiralling cost that could offset any savings gained from having a systematic resource allocation framework. To the best of our knowledge, there is limited work focused on making deep learning pipelines practical for companies managing data lakes at similar scale.

Our work is directed towards improving the cost efficiency for deep learning model training over a query's logical plan. We focus on optimizing the state-of-the-art tree convolution (Tree CNN) based pipelines. Present pipelines [20, 21] apply Tree CNN over full query plans and have poor per-batch memory footprint and long epoch runtime. In turn, we propose *Prestroid*, a pipeline that addresses these problems and is cheaper for training. We denote our contributions below.

- We present a case study over SQL queries issued over Grab's data lake. Our study revealed a large number of distinct query predicates and a wide disparity in query plan sizes, which are significantly different to TPC benchmarks. We believe that these features can be common in other large-scale workloads. We show why direct applications of present deep learning pipelines are 1) not space-efficient for encoding and 2) creates excessive padding that mandates the use of higher tier and more expensive GPUs for training.
- We present the components of *Prestroid*, consisting of a sub-tree convolution model, a Word2Vec model for controlling plan level embedding and a novel sub-sampling algorithm for decomposition of large query plans whilst preserving breadth level information for tree convolution. *Prestroid* is cost-efficient as it achieves better forecast accuracy at lower per-batch memory footprint and faster epoch runtime.
- We demonstrated experimentally that our pipeline enables model training cost reduction of 2x and 13.2x at batch sizes of 32 and 256 respectively. Moreover, *Prestroid* achieves better predictive accuracy than state-of-the-art, which helps the resource provisioning for large-scale workloads in Grab.
- We publicly avail our Grab-Traces and TPC-DS dataset¹. To the best of our knowledge, Grab-Traces is the largest available industry-based dataset of query plans. Our technical report may also be found at [14].

2 BACKGROUND & RELATED WORK

2.1 Cloud based resource utilization

Majority of cloud platforms offer pay-as-you-use resource type customization to meet the varied demands of customers [35]. These resources are packaged as a tiered set of on-demand virtual machines (VMs) with different cores, memory & storage capacities and pricing. A collection of VMs forms a *cluster* and are the workhorses of big data query engines [5, 33] on the cloud. As it is possible to add an indefinite number of VMs to a cluster (termed as *scale out*),

developing a good resource forecasting framework for projected query workload would enable the selection [13, 31] of just the right combination of VMs to meet demands at cost optimal pricing.

2.2 State-of-the-art

Analytical models [12, 30, 32, 39] attempt to shed light on the internal mechanisms of a query engine by modelling various aspects, such as data access and workload scheduling. Such models have been developed extensively for transactional based systems such as MySQL or distributed processing engines such as Hive and Hadoop. However, such models are highly specific to a single engine and are hard to develop.

Machine learning models approach the problem differently. Earlier models explored simple techniques, such as KCCA [9], for subspace mapping of queries-resources, or standard regression analysis [29]. In recent years, there has been a shift towards deep learning models such as feed forward network [17, 26], Tree CNN [21], RNN based networks [27, 34] or reinforcement learning [20, 22]. These models were designed to capture the inherent complexities of OLAP queries to which simpler models failed to do so.

2.3 Query feature extraction

Deep learning models are trained and evaluated over numerical inputs. This implies the need to formulate methods that translate plain SQL into their vectorized representations.

SQL parsing - A simple approach would be to parse a query string entirely to aggregate key features that represents the query. Gnapathi et al. [9] proposed 9 distinct features that characterizes an SQL query, where as Makiyama et al. [19] suggested representing a query as a collection of the weighted frequencies of its individual word tokens.

Logical plan parsing - To venture further, aggregations may be done over a query's logical plan structure. A logical plan is directed acyclic graph (DAG) representation of the operations needed to be fulfilled before the final table can be materialized. Each operation is represented as a node in the DAG. Logical plans allow deeper insights into the execution sequence taken by the query engine that may not be attained through plain text parsing. Such were the works proposed by [2, 3, 9], in which the authors represented a query as vector aggregation of specific operations within a plan. In modern database engines such as Presto or MySQL, obtaining a query's logical plan can be easily achieved using the "*EXPLAIN <text>*" key word without the need to execute the query.

O-T-P encoding - The *Operator-Table-Predicate* encoding format has been adopted by many state-of-the-art work in the field of deep learning for query-resource forecasting [17, 21, 27]. At its core, a distinction is made between the categories of *Operators*, which are wildcards representing key operations such as joins or projections, *TABLES* indicating the scanned tables and *PREDICATES* indicating the conditions over which data is filtered. Different encoding techniques may be applied within each category and the resultant combination of {O,T,P} is the feature representation of that query.

Such encoding techniques (with slight variants) were adopted in [17, 21, 27, 34] and worked well for a good variety of models, from simple feed forward networks [17], recurrent neural networks [27, 34] and Tree CNN networks [21]. We adopt the O-T-P encoding

¹Dataset github: <https://github.com/grab/grab-query-traces>

approach in this work by first casting a query into its logical plan, before re-casting into a binary tree comprising only of O/T/P nodes.

2.4 Tree CNN based models

SQL featurization, as a standalone, may fall short of adequately representing a query as they fail to account for the order of executions captured within the plan sequence [25]. In the context of database engines, the choice of plan may significantly impact run time performance [4, 20, 28]. Such sequence order sensitivity implies the need to develop models that are able to differentiate plan level spatial arrangements, in order to maximize predictive accuracy.

Tree CNNs [21] are one such model. They draw inspiration from pre-existing convolution networks applied over images or graphs. The pioneering works of Mou et al. [25] have inspired the application of Tree CNN based models in the field of query plan selection [21] and natural language classification [6, 8]. Tree CNN networks aggregate information between parent and children by sliding and pooling triangular kernels breadth first across each node, thereby capturing the positional ordering of operators within a plan. The reader is encouraged to review [25] for a deeper understanding.

3 THE CURSE OF DIVERSITY & SCALE

3.1 An industry case study

Figure 2: Contrast of 245,849 logical plan samples from Grab alongside 103 TPC-DS & 22 TPC-H publicly available samples. We showed theoretical plots for skewed trees and balanced binary trees for reference.

Grab is Southeast Asia’s leading superapp that provides everyday services such as mobility, deliveries (food, packages, groceries), mobile payments and financial services to millions of Southeast Asians. Rapid expansion across the region has resulted in queries that are vastly different in characteristics and resource needs. We plotted a sample of 245,849 logical plans, obtained over 2 months, on their node count and maximum depth in Fig 2. Here, the maximum depth refers to the largest distance between the root and any leaf node. As reference, we contrasted them with theoretical plots for balanced binary and skewed trees (left-deep trees with only 1 child). An observation was that majority of the plans straddled in between both plots, indicative of large plan diversity.

Distinction from public datasets - We contrast the plans from Grab with publicly available TPC-DS & TPC-H plans in Fig 2. *Notably, these plans covered a smaller range of distribution relative to the plans from Grab, implying that the latter was richer in quantity and the span of plan sizes.* To quantify, the maximum plan (size, depth) observed was (477, 38) for TPC-H, (883, 73) for TPC-DS and (4969, 321) for Grab.

This distinction highlights an important area of research that has yet to be explored. To the best of our knowledge, most research fail to address the issues that surface in training Tree CNN based deep learning pipelines over a large and diverse set of logical plans. We attribute one possibility to the fact that publicly available query patterns lack both the quantity and permutations needed to replicate the scenario present in Grab. Addressing these issues is a critical step forward in bridging the gap between research and applications to the industry.

Dynamism of query patterns - The performance of a deep learning model is highly dependent on the training dataset. As such, training frequency should be contingent on how rapidly new data is introduced. In order to quantify this rate, we sampled 373K Presto queries from Grab across 1 month span and extracted all tables required by the queries. We asked ourselves if the model was used to predict for a subsequent window of W days, what is the percentage of tables in the new queries that the model has not encountered.

W	1	3	5	7	9
%	1.65	4.76	7.64	9.27	12.18

Table 1: Percentage of new tables that a model has not seen over the next W days window.

We are observing a high rate of growth in tables within the company as seen in Table 1. This motivates the frequent re-training of our models. For example, at $W = 9$, the model has been used to predict new queries over the next 9 days. In doing so, the model suffers from a high degree of inaccuracy, given that 12.18% of tables scanned are new tables that the model has not been trained over. We therefore recommend the daily re-training of our models in practice.

3.2 0-padding for dimensional consistency

It is common practice to implement *NULL*, or “0-padding” in order to reconcile irregularities in input data dimensions. For Grab, these irregularities appear due to having both large and small query plans in the training data. Theoretically, 0-padding will not impact model training performance as a null input does not affect weight updates. Yet 0-padding introduces redundant information in the model [38] that has consequences.

For a given training batch size, excessive 0-padding will lead to an increase in overall per-batch memory footprint. This leads to longer data transfer time between CPU-GPU for each epoch cycle and more computations needed over the data, resulting in longer per-epoch runtime. More importantly, scarce GPU memory bandwidth will be exhausted for models with multiple layers as the GPU has to retain all intermediate data to compute back propagation gradients. To Grab, the implications are twofolds. Machine learning practitioners either do not have the flexibility to tune models over a wide range

of batch sizes, which may lead to sub-optimal model performances [15], or they have to scale out their hardware to more expensive GPU tiers on the cloud, which is a cost concern.

One technique to avoid unnecessary 0-padding is down sampling. This problem has been largely explored in the field of image processing [23, 36]. Unfortunately, such cannot be said for the field of deep learning over tree based structures.

3.3 Surge in query predicates

A query predicate defines a condition to be applied over transformations within each stage of a query plan. In the case of conditional filters, predicates are represented as a set consisting of {Columns, Comparison operator, Filter values} [34].

Our analysis of the query patterns in Grab revealed that while the number of tables being queried are few, the number of unique predicates were very large relative to publicly available datasets². This is understandable, given that a single table can have multiple columns for performing filters or joins and that predicates may vary in complexities according to business rules. In turn, we highlight the flaws of existing encoding techniques applied to these predicates.

1-Hot [17] - This may cause a sizeable increase in the encoded vector's length, creating sparse vectors of a single 1 and remaining 0s, which occupies a large chunk of space.

Value normalization [37] - Normalization is used to constrain filter values to a (0,1) range suitable for training. This technique works mostly for integer and floats. For strings, dictionary encoding may be used to cast it to an integer. Unfortunately, this technique do not work for predicate columns and must be coupled with others as discussed.

R-vectors [21] - The R-vector representation was proposed by Marcus et al. in an attempt to capture the semantic relations between column values in a database. A Word2Vec model is trained, for encoding newly materialized tables for each query, by first treating each row as a sentence and each column as a token. While R-vectors enable compact feature representations of queries, they are costly to deploy. Prior execution of each query is needed to materialize the table for encoding. This clearly does not work well for hundreds of thousands of queries, where each may take hours to complete.

3.4 Summary

In summary, our observations of the scale and diversity of query workloads in Grab surfaces several problems with existing deep learning pipelines. Firstly, the need to reconcile both large and small query plans begets excessive 0-padding. Secondly, present encoding techniques for large query dataset are space inefficient. These factors amount to an increase in batch memory footprint and induces unnecessary computations. Consequently, model training has to be done over longer horizons and on scaled out GPU machines, which are costly overall. Directly addressing these problems will yield improvements to the cost efficiencies of such pipelines.

²Based on 19,876 sampled queries from Grab in our training dataset, the number of unique predicate counts may extend as much as 30,707. In contrast, 5,153 TPC-DS queries, generated from 81 templates, yield a count of 1,450.

4 DESIGN AND IMPLEMENTATION

Here we present Prestroid's data pipeline for addressing the challenges in Section 3. We present an overview of our end-to-end model training process before explaining how we reduced batch memory footprint through minimizing node level encoding and 0-padding. For simplicity, we focus on single objective learning in which the model has to predict the *total CPU time* of a query.

4.1 Building the data pipeline

Prestroid consists of a data pre-processing phase and a model training phase, illustrated in Fig 3. Firstly, an incoming query is decomposed into its logical plan and further re-cast using the *O-T-P* framework. We apply the following rules:

- For a non-join node N_I , set its type to be an OPR. Create a right child as type PRED with the predicate value. The left child is untouched.
- For a join node N_J , set its type to be an OPR. The left & right children are untouched.
- For a leaf node N_L (table scans), set its type to be an OPR. Create a left child as type TBL with the table scanned as value. Create a right child as $\emptyset$.
- Transform the resultant tree into a binary tree by adding $\emptyset$ to any node with fewer than 2 children.

Plan encoding - OPR & TBL nodes are collected separately and 1-Hot encoded whilst PRED nodes are encoded using our model in Section 4.2. We traverse each node in a tree and apply respective encoding in the [OPR, PRED, TBL] format.

Sub-tree model - We apply our sub-sampling algorithm and select the first K sub-trees as representative features for a query. Our model uses 3 layers of CNN. We then apply bit masking and perform one-way dynamic pooling [25] over each sub-tree. Finally, we flatten all sub-trees into a single vector, pass them through 2 dense layers with ReLU activation before a single layer with sigmoid activation as prediction.

4.2 Learned predicate embedding

Word2Vec model - The goal here is to identify an n-dimensional feature space that enables control over each predicate embedding in a meaningful way. For example, the words "LONGITUDE" and "LATITUDE" appear frequently with each other in the queries that we sampled. We would expect them to be spatially closer in our feature space as compared to "DATAMART", which is used in a totally separate context.

Such a problem has been explored extensively in the arena of natural language. Popular models proposed are Word2Vec models [24] such as skip-gram or CBOW. We show that such models can be used for learning predicate representations based on logical plan extracts from queries. The key idea is to train our Word2Vec models over all predicate tokens with values omitted. To illustrate, consider the example in Fig 4. To train our model, all conjunctions and values from each predicate are first stripped off, leaving behind only the columns and comparison operators. We then train our token sets using the Word2Vec model offered by Python's Gensim package. We ran the model using a window size of 5, minimum token count of 10 and a range of feature sizes. Tuning the feature size allows us to control the encoding space for predicates.

Figure 3: Prestroid data pipeline used for model training & prediction

Handling conjunctions - After obtaining our trained encoding, we cast our predicate into a tree where the nodes are either conjunctions (AND or OR) or a single predicate clause. For the latter, we encode each word token and take the overall average as the node level encoding. We then apply MIN feature pooling over all children nodes for AND conjunctions and MAX feature pooling for OR conjunctions, following [34].

Out of vocabulary tokens - One approach to address out of vocabulary tokens encountered during pipeline deployment is to follow a hierarchy of updates. For example, in order to encode unseen predicates, we first search for PRED nodes within the query and take their average features. If not, we take the average of all tokens in the query. If all else fails, we take the average encoding for all PRED nodes in the global set. This simple approach works well in our experiments.

Figure 4: Illustration of our Word2Vec model training and encoding over a toy predicate example

4.3 Sub-tree sampling

The goal here is to decompose a large query tree into smaller sub-trees. Thereafter, by careful selection of the top K representative sub-trees, we are able to reduce the overall 0-padding needed to our model. Algorithm 1 denotes the pseudo code for our sub-tree sampling algorithm. In contrast to naive breadth first or depth

first pruning, our sub-sampling algorithm ensures that information needed during Tree CNN is preserved.

The crux of the algorithm lies in the observation that, given a root node **R**, node count limit **N** and **C** layers of convolution, **R** only requires information present in its children up till **C** levels below. The rough workings of our algorithm are as such.

- Starting from **R** (depth 0), we consider the possibility growing our sub-tree at incremental depths of 1. At each depth, all children must be materialized.
- Where the sub-tree has size exceeding **N** at depth **D**, we regress back to depth **D-1** and prune the tree breadth first.
- All nodes up till depth $(D-C-1)$ have complete information and are allowed to vote³. We set their votes to 1. Nodes beyond depth $(D-C-1)$ have their votes set to 0.
- The algorithm is repeated for all nodes at depth $(D-C)$.

Our algorithm reduces the use of 0-padding by enabling the user to tune the values of **N** / **K**, which reduces the input Tensor dimensions to the model. Current implementation treats each plan as a binary tree and enforce that $N > 2^{C+1} - 1$.

5 EXPERIMENTAL EVALUATION

Here we designed 3 experiments to evaluate Prestroid relative to selected benchmarks. *Exp 1* - We assess how well Prestroid performed in terms of MSE score. All models were trained for 3 rounds, with average MSE scores taken from the best performing iterations. In each round, early stopping was employed to prevent model over fitting. *Exp 2* - We evaluate the accuracy of Prestroid in forecasting suitable resources quantities. *Exp 3* - We evaluate the per-batch memory footprint, required epoch training time and training costs for Prestroid sub-tree models over cloud based infrastructures.

³ A vote is simply a bit masked value that is applied after all convolution layers the model. A vote of 1 is assigned only to nodes that have complete information and can be used as valid signals during post-convolution.

Algorithm 1: Sub-tree sampling algorithm

input: N node limit, C convolution layers, R root node
output: $[S_n]$ sub-sample binary trees, $[V_n]$ - Votes
constraints: $N > 2^{C+1} - 1$

def *getNode*(R: root node, D: depth):
 returns all nodes until depth D, including R

$\emptyset \leftarrow$ FIFO queue Q_f
 $[\] \leftarrow S, V$
 $Q_f.enqueue(R)$

while not $Q_f.empty$ **do**
 node = $Q_f.pop()$
 candidates = $[\]$
 depth = 0

 # Terminate when we hit leaf nodes or limit
while $len(candidates) \leq N$ **do**
 prior_candidates = candidates
 depth += 1
 candidates = *getNode*(node, depth)
 if $len(candidates) == len(prior_candidates)$ **then**
 # No new children nodes to discover
 break

end

 sub_tree = prior_candidates
 sub_tree_count = $len(sub_tree)$
if $len(candidates) == len(prior_candidates)$ **then**
 # All nodes are valid for complete tree
 votes = $[1] * sub_tree_count$

else
 eligible = $len(getNode(node, depth - C - 1))$
 vote_nodes = $[1] * eligible + [0] * (sub_tree_count - eligible)$
 $Q_f.enqueue(getNode(node, depth - C))$
 S.append(sub_tree)
 V.append(votes)

end
 return (S, V)

5.1 Experimental Setup

Infrastructure - Our models were trained using TensorFlow [1] over Azure. We used the NC_V3 series powered by NVIDIA Tesla V100 GPUs with 16GB memory. Exp 1 & 2 was conducted on NC12s_V3 cluster with 2 GPUs. Exp 3 was conducted over NC6s_V3 / NC12s_V3 / NC24s_V3 clusters fully utilizing all 1 / 2 / 4 GPUs.

Dataset - Our dataset was based on 2 different sources.

TPC-DS: We generated a total of 5,153 unique queries with 81 unique templates from the TPC-DS Hive dataset. We filtered all queries with total CPU time between 1 - 60 min. All queries were executed on Presto at a scale factor of 10. We applied log transformation followed by min-max normalization over all recorded total CPU time to constrain all training values in between 0 - 1. We used a split ratio of 8 / 1 / 1 for training / validation / testing. Splitting was done at the template level.

Grab-Traces: We curated a 2 month sample of query traces from Grab. These queries were executed across multiple Presto clusters in deployment. Only successfully executed queries were selected. We first filtered all queries with total CPU time between 1 - 60 min. We then applied log transformation followed by min-max normalization to constrain all training values in between 0 - 1. The resultant dataset contained 19,876 queries split into 8 / 1 / 1 ratio for training / validation / testing.

Comparisons - All deep learning model comparisons were trained using ADAM [16] optimizer, batch size of 64 and the Huber loss function, unless stated otherwise.

- **Log binning** [7] - We split all query plans by their node counts into B log bins. The average total CPU timing is taken within each bin and used for inference. We used this as a naive benchmark for comparison with other models. Experimentally, we found that the optimal values for B were 1000 & 20 for Grab-Traces & TPC-DS dataset respectively.
- **SVR** [9] - A support vector regression (SVR) model is trained using direct query parsing and plan operator instance counts. We omitted plan operator cardinalities as part of the feature vector. We found that the best performing models used a polynomial kernel of degree 4 and sigmoid kernel of degree 3 for Grab-Traces & TPC-DS dataset respectively.
- **Modified MSCN** [17] - We modified the multi-set convolutional network (M-MSCN) for our task. Although MSCN was built for cardinality estimation, its design principles were based on Deep Sets [40], which we argue allows the network to be generalized for query-cost estimation. We used 0-padding for input consistency and set dropout to 5%. We set learning rate of $1e^{-3}$ and 256 units per layer for Grab-Traces and $1e^{-4}$ and 24 for TPC-DS.
- **WCNN** [41] - We followed the Word Convolution (WCNN) implementation and optimal hyper-parameters as reported. We explored the use of {100, 250} kernels for each of the {3, 4, 5} sliding window convolution filters. We used a token embedding layer of dimensions 100 in our network. Dropout was set at 50%, batch size at 16 and learning rate at $1e^{-3}$ for Grab-Traces and $1e^{-4}$ for TPC-DS.
- **Prestroid (Full- P_f)** - To show the gains from using sub-trees, we implemented Prestroid over full query plans without any tree pruning. This model is similar to the tree convolution segment of [21]. Here, P_f represents the feature size chosen from our Word2Vec model. We explored the range of $P_f \in \{100, 200, 300\}$ and $\{50, 100\}$ for Grab-Traces & TPC-DS dataset respectively. We set learning rate of $1e^{-4}$.

5.2 Hyper-parameter tuning

Prestroid exposes 3 new parameters that may be tuned for performance. They are 1) $P_f \rightarrow$ predicate features size, 2) $K \rightarrow$ number of sub-trees chosen to represent a query and 3) $N \rightarrow$ max node count per sub-tree. We explored 2 variations of $N \in \{15, 32\}$. For Grab-Traces, we explored $P_f \in \{100, 200, 300\}$, $K \in \{5, 9, 21\}$ where $N = 15$ and $K \in \{5, 11, 20\}$ where $N = 32$. For TPC-DS, we explored $P_f \in \{50, 100\}$, $K \in \{31, 43, 47\}$ where $N = 15$ and $K \in \{20, 28, 32\}$ where $N = 32$.

We used 512 / 512 / 512 CNN kernels and 128 / 64 perceptron units for Grab-Traces. For TPC-DS, we scaled the architecture down

Models	Epoch	MSE
Log bins	-	96.91
SVR	-	106.16
M-MSCN	78	66.35
WCNN-100	55	50.35
WCNN-250	55	50.90
Full-100	52	50.82
Full-300	51	48.16
Prestroid (15-9-300)	49	49.23
Prestroid (32-11-200)	41	46.09

(a) Performances on Grab-Traces dataset

Models	Epochs	MSE
Log bins	-	58.09
SVR	-	58.97
M-MSCN	17	145.91
WCNN-100	15	100.62
WCNN-250	29	103.05
Full-50	75	58.33
Full-100	69	55.60
Prestroid (15-47-50)	46	46.61
Prestroid (32-32-100)	49	47.24

(b) Performances on TPC-DS dataset

Table 2: Recorded MSE errors ($minutes^2$) for best performing Prestroid sub-tree (N-K- P_f) models, Prestroid full tree models and respective comparisons. We also included the highest observed epoch at convergence out of all 3 runs.

to 128 / 128 / 128 and 32 / 8. We set dropout as 10% for kernel and bias weights, with batch normalization in between each dense layer. We used ADAM with learning rate of $1e^{-4}$ and optimized for Huber loss. Herein, We refer to any future variations of Prestroid as Prestroid (N-K- P_f).

Performance on Grab-Traces - Our results on the Grab-Traces dataset imply that sub-tree models have greater learning capacities than other state-of-the-art deep learning models. The optimal sub-tree configuration observed was (32-11-200).

Compared to full tree models, our sub-tree models have K times more features as inputs to the dense layer after convolution, unlike the former which collapses all plan level features into a single vector via dynamic pooling. Scaling up the inputs by K times enabled our sub-tree models to learn a richer set of mappings between query plans and cost estimates.

Surprisingly, WCNN showed comparable performances to Prestroid Full-100, implying that convolution models that operate directly on SQL strings are able to extract just as much information as compared full query plans. However, SQL strings do not reflect the true cost of how a query is executed. Take the case of a simple command "SELECT * FROM A, B, C". Whilst WCNN understands that multiple joins are performed for tables A, B & C, the join ordering and the type of join used is hidden from the model. These details are decided by the query optimizer at runtime and can only be accessed at the plan level. Consequently, this limits the learning ability of WCNN as compared to Prestroid sub-trees, which we showed had better learning than full trees.

Finally, both Prestroid full and sub-trees generally fared better than SVR and M-MSCN. The latter models were trained over an aggregation of features from both query and logical plans, which caused valuable plan level information to be lost. In contrast, Prestroid was able to leverage these signals through the use of triangular kernels tuned to detect the spatial patterns between parent and children nodes.

Performance on TPC-DS - While we observed that Prestroid (15-47-50) yielded the lowest MSE score over the TPC-DS dataset, the shortage of query variations yields two interesting observations. The first is that simpler models (Log Binning & SVR) showed comparable performance relative to deep learning models, a trend absent when benchmarked over Grab-Traces. We assert that the latter models are harder to train and require a broad range of query characteristics and training data size to outperform our naive baselines.

Unfortunately, such features were absent in the TPC-DS dataset, which contains only 103 publicly available query templates.

The second is a sharp decline in the performance for WCNN. WCNN models are relatively heavy; WCNN-100 contains 363,301 trainable parameters whereas Full-100 contains 195,469 trainable parameters. This implies that it is easier for WCNN to overfit when limited training data is present.

5.3 Resource allocation

Figure 5: Percentage of cluster resources that were over / under allocated for Grab sampled query workloads. The lower the magnitude, the better the accuracy.

We designed an experiment to evaluate the resource allocation accuracy of Prestroid. Our results in Fig 5 were based off a test data set of 1,987 query traces from Grab-Traces. We categorized our results into 2 groups: *Over provisioned* & *Under provisioned*, with the intent of understanding by how much percentage of actual cluster resources did our sub-tree models over/under allocate to execute these queries. For example, queries which our model assigned excess CPU time were classified as Over provisioned. We compared Prestroid (15-9-300) & (32-11-200), which were the best overall performers, against Prestroid (Full-100) & (Full-300). Our observations were that Prestroid sub-trees tend to perform better for both over allocated and under allocated queries. Overall, all models generally under provisioned resources (see yellow bar), with Prestroid sub-trees achieving better resource allocation accuracy. This highlights the importance of our pipeline in enabling Grab to achieve optimal cost strategies for resource allocation.

5.4 Optimizing hardware usage

We studied how our encoding and sub-sampling approach improves batch efficiency and training speeds relative to full-tree models in Fig 6. In turn, these improvements may be translated to cost savings over cloud resources.

Figure 6: Top: Average per-batch memory footprint (MB) for various models at batch size of 32 over Grab-Traces dataset. *Pr* to denote Prestroid models. Bottom: Average epoch time (minutes) for various models.

Reduction in batch size - The impact of 0-padding was significant given input size irregularities in Grab-Traces query plans. For a batch size of 32, comparisons between Prestroid (15-9-300) & (32-11-200) relative to Prestroid (Full-300) yields a reduction of 13.5x and 5.8x respectively. 0-padding was used to enforce dimensional consistency of inputs to all models. Padding for full tree models was based on the size of the largest tree (1945 nodes). This meant padding a skewed distribution of small trees to the size of the largest.

Reduction in training time - Reductions in input data size consequently leads to faster training times owing to improved bandwidth utilization for data transfer between CPU to GPU and fewer computations over the data. Comparisons between Prestroid (15-9-300) & (32-11-200) relative to Prestroid (Full-300) show speedups of 3.45x and 2.6x. However, we also observed a disproportionate growth in epoch runtime for larger selections of K. We attribute this to an inefficiency in our current code, which uses Tensorflow's *tf_map* operator to perform sequential convolutions over each sub-tree. This limitation may be addressed in future work.

Other comparisons - For completeness, we have included the per-epoch run time and batch sizes for M-MSCN and WCNN. For M-MSCN, a large number of distinct predicates coupled with variations in table, join and predicate sets per query produced sparse and large input tensors to the model, creating long epoch run times and high memory footprint. In contrast, the use of a trainable embedding layer in WCNN allowed us to minimize inputs to a single 1-D vector. This is highly efficient for speed and memory footprint reduction. Yet WCNN has shown sub-optimal performances relative to Prestroid sub-tree models in Table 2, which we argue is of first importance when designing query-cost estimation models.

Cost Savings - Finally, we evaluated the cost of model training using Azure's NC6s_V3 / NC12s_V3 / NC24s_V3 clusters. We chose the lowest possible cost among all clusters that permitted training with a specified batch size. Each cluster contained 1 / 2 / 4 GPUs respectively and were priced per hour⁴ at \$4.23 / 8.47 / 18.63. In multi-GPU clusters, data parallelism [18] was employed to distribute batch workloads. We compared Prestroid (15-9-300) & (32-11-200), which were the best performers, against Prestroid (Full-100) & (Full-300). We assumed model training until the epochs denoted in Table 2a.

Batch size costs - We observed diminishing cost returns due to communication overheads and a non-linear increase in pricing for multi-GPU clusters. Consequently, it is economically cheaper to train using a single GPU. Fig 7 suggests that our sub-tree models were cheaper to train across a range of batch sizes as they were faster and had lower memory footprint. For large batches, Full tree models had to be trained on multi-GPUs due to out of memory errors, while sub-tree models could still be trained on a single GPU. To quantify the impact of our sub-sampling approach, we observed training cost decline of \$76.25 to \$5.79 in switching from Prestroid (Full-300) to Prestroid (15-9-300) for a batch size of 256.

Figure 7: Lower bound of costs using Azure NC_V3 clusters for training with varying batch sizes. P-15* & P-32* denotes Prestroid (15-9-300) & (32-11-200) respectively

6 CONCLUSION

This paper tackles the challenge of extending deep learning Tree CNN models to large and diverse queries in a cost efficient manner. Our sub-tree model exposes three levers which users can tune to control model accuracy, batch size and epoch training time. Careful selection of these parameters allows one to accelerate the training process over tens of thousands of query plans at lower memory and high accuracy. This enables model training to be done faster and over cheaper cloud resources, leading to substantial cost savings. Although our experiments were conducted using Grab specific data, we believe that the general techniques can be distilled and applied to any other company managing data lakes at a similar scale.

7 ACKNOWLEDGEMENTS

This work was funded by the Grab-NUS AI Lab, a joint collaboration between GrabTaxi Holdings Pte. Ltd. and National University of Singapore. We thank See-Kiong Ng, Hannes Kruppa and Rahul Pentti for their support and advice.

⁴Rates were up-to-date as of the time of writing and may change in future

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*. 265–283.
- [2] Mert Akdere, Ugur Cetintemel, Matteo Riondato, Eli Upfal, and Stanley B Zdonik. 2011. The Case for Predictive Database Systems: Opportunities and Challenges.. In *CIDR*, Vol. 2011. 167–174.
- [3] Mert Akdere, Ugur Cetintemel, Matteo Riondato, Eli Upfal, and Stanley B Zdonik. 2012. Learning-based query performance modeling and prediction. In *2012 IEEE 28th International Conference on Data Engineering*. IEEE, 390–401.
- [4] Michael Armbrust, Reynold S Xin, Cheng Lian, Yin Huai, Davies Liu, Joseph K Bradley, Xiangrui Meng, Tomer Kaftan, Michael J Franklin, Ali Ghodsi, et al. 2015. Spark sql: Relational data processing in spark. In *Proceedings of the 2015 ACM SIGMOD international conference on management of data*. 1383–1394.
- [5] Ekaba Bisong. 2019. Google BigQuery. In *Building Machine Learning and Deep Learning Models on Google Cloud Platform*. Springer, 485–517.
- [6] Nghi DQ Bui, Lingxiao Jiang, and Yijun Yu. 2017. Cross-language learning for program classification using bilateral tree-based convolutional neural networks. *arXiv preprint arXiv:1710.06159* (2017).
- [7] James Dougherty, Ron Kohavi, and Mehran Sahami. 1995. Supervised and unsupervised discretization of continuous features. In *Machine learning proceedings 1995*. Elsevier, 194–202.
- [8] Bui Nghi DQ, Yijun Yu, and Lingxiao Jiang. 2019. Bilateral dependency neural networks for cross-language algorithm classification. In *2019 IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 422–433.
- [9] 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 *2009 IEEE 25th International Conference on Data Engineering*. IEEE, 592–603.
- [10] Chetan Gupta, Abhay Mehta, and Umeshwar Dayal. 2008. PQR: Predicting query execution times for autonomous workload management. In *2008 International Conference on Autonomic Computing*. IEEE, 13–22.
- [11] Bingsheng He, Mao Yang, Zhenyu Guo, Rishan Chen, Bing Su, Wei Lin, and Lidong Zhou. [n.d.]. Comet: batched stream processing for data intensive distributed computing. In *Proceedings of the 1st ACM symposium on Cloud computing*. 63–74.
- [12] Jiamin Huang, Barzan Mozafari, Grant Schoenebeck, and Thomas F Wenisch. 2017. A top-down approach to achieving performance predictability in database systems. In *Proceedings of the 2017 ACM International Conference on Management of Data*. 745–758.
- [13] Sangeetha Abdu Jyothi, Carlo Curino, Ishai Menache, Shravan Matthur Narayana-murthy, Alexey Tumanov, Jonathan Yaniv, Ruslan Mavlyutov, Íñigo Goiri, Subru Krishnan, Janardhan Kulkarni, et al. 2016. Morpheus: Towards automated slos for enterprise clusters. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*. 117–134.
- [14] Johan Kok Zhi Kang, Gaurav, Sien Yi Tan, Feng Cheng, Shixuan Sun, and Bingsheng He. 2021. Efficient Deep Learning Pipelines for Accurate Cost Estimations Over Large Scale Query Workload. *arXiv eprint arXiv:2103.12465* (2021).
- [15] Nitish Shirish Keskar, Jorge Nocedal, Ping Tak Peter Tang, Dheevatsa Mudigere, and Mikhail Smelyanskiy. 2019. On large-batch training for deep learning: Generalization gap and sharp minima. In *5th International Conference on Learning Representations, ICLR 2017*.
- [16] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [17] Andreas Kipf, Thomas Kipf, Bernhard Radke, Viktor Leis, Peter Boncz, and Alfons Kemper. 2019. Learned cardinalities: Estimating correlated joins with deep learning. In *CIDR*.
- [18] Alex Krizhevsky. 2014. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997* (2014).
- [19] Vitor Hirota Makiyama, Jordan Raddick, and Rafael DC Santos. 2015. Text Mining Applied to SQL Queries: A Case Study for the SDSS SkyServer.. In *SIMBig*. 66–72.
- [20] Ryan Marcus, Parimarjan Negi, Hongzi Mao, Nesime Tatbul, Mohammad Alizadeh, and Tim Kraska. 2020. Bao: Learning to Steer Query Optimizers. *arXiv preprint arXiv:2004.03814* (2020).
- [21] Ryan Marcus, Parimarjan Negi, Hongzi Mao, Chi Zhang, Mohammad Alizadeh, Tim Kraska, Olga Papaemmanouil, and Nesime Tatbul. 2019. Neo: A Learned Query Optimizer. *Proc. VLDB Endow.* 12, 11 (July 2019), 1705–1718. <https://doi.org/10.14478/3342263.3342644>
- [22] Ryan Marcus and Olga Papaemmanouil. 2019. Towards a Hands-Free Query Optimizer through Deep Learning.. In *CIDR*.
- [23] Dmitrii Marin, Zijian He, Peter Vajda, Priyam Chatterjee, Sam Tsai, Fei Yang, and Yuri Boykov. 2019. Efficient segmentation: Learning downsampling near semantic boundaries. In *Proceedings of the IEEE International Conference on Computer Vision*. 2131–2141.
- [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] Lili Mou, Ge Li, Lu Zhang, Tao Wang, and Zhi Jin. 2016. Convolutional Neural Networks over Tree Structures for Programming Language Processing. In *AAAI*.
- [26] Jennifer Ortiz, Magdalena Balazinska, Johannes Gehrke, and S Sathya Keerthi. 2018. Learning state representations for query optimization with deep reinforcement learning. In *Proceedings of the Second Workshop on Data Management for End-To-End Machine Learning*. 1–4.
- [27] Jennifer Ortiz, Magdalena Balazinska, Johannes Gehrke, and S Sathya Keerthi. 2019. An empirical analysis of deep learning for cardinality estimation. *arXiv preprint arXiv:1905.06425* (2019).
- [28] Yongjoo Park, Shucheng Zhong, and Barzan Mozafari. 2020. Quicksel: Quick selectivity learning with mixture models. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 1017–1033.
- [29] Adrian Daniel Popescu, Vuk Ercegovic, Andrey Balmin, Miguel Branco, and Anastasia Ailamaki. 2012. Same queries, different data: Can we predict runtime performance?. In *2012 IEEE 28th International Conference on Data Engineering Workshops*. IEEE, 275–280.
- [30] Kaushik Rajan, Dharmesh Kakadia, Carlo Curino, and Subru Krishnan. 2016. Perforator: eloquent performance models for resource optimization. In *Proceedings of the Seventh ACM Symposium on Cloud Computing*. 415–427.
- [31] Rathijit Sen, Alekh Jindal, Hiren Patel, and Shi Qiao. 2020. AutoToken: predicting peak parallelism for big data analytics at Microsoft. *Proceedings of the VLDB Endowment* 13, 12 (2020), 3326–3339.
- [32] Rathijit Sen and Karthik Ramachandra. 2018. Characterizing resource sensitivity of database workloads. In *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 657–669.
- [33] Raghav Sethi, Martin Traverso, Dain Sundstrom, David Phillips, Wenlei Xie, Yutian Sun, Nezih Yegitbasi, Haozhun Jin, Eric Hwang, Nileema Shingte, et al. 2019. Presto: Sql on everything. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE, 1802–1813.
- [34] Ji Sun and Guoliang Li. 2019. An end-to-end learning-based cost estimator. *arXiv preprint arXiv:1906.02560* (2019).
- [35] Luis M Vaquero, Luis Rodero-Merino, and Rajkumar Buyya. 2011. Dynamically scaling applications in the cloud. *ACM SIGCOMM Computer Communication Review* 41, 1 (2011), 45–52.
- [36] Thang Vu, Cao Van Nguyen, Trung X. Pham, Tung M. Luu, and Chang D. Yoo. 2018. Fast and Efficient Image Quality Enhancement via Desubpixel Convolutional Neural Networks. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*.
- [37] Lucas Woltmann, Claudio Hartmann, Maik Thiele, Dirk Habich, and Wolfgang Lehner. 2019. Cardinality estimation with local deep learning models. In *Proceedings of the Second International Workshop on Exploiting Artificial Intelligence Techniques for Data Management*. 1–8.
- [38] Sai Wu, Mengdan Zhang, Gang Chen, and Ke Chen. 2017. A new approach to compute cnns for extremely large images. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 39–48.
- [39] Wentao Wu, Yun Chi, Shenghuo Zhu, Junichi Tatemura, Hakan Hacigümüs, and Jeffrey F Naughton. 2013. Predicting query execution time: Are optimizer cost models really unusable?. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*. IEEE, 1081–1092.
- [40] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabás Póczos, Ruslan R Salakhutdinov, and Alexander J Smola. 2017. Deep Sets. In *NIPS*.
- [41] 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*. 209–224.