

A Recommendation System for Execution Plans

*Ms. Pratibha V. Waje, #Dr. R. Jain

*Assistant Professor, SVIT, Nashik, Maharashtra, India.

#Professor, CSJMU, Kanpur, India.

Abstract: In this paper, we discuss an approach for execution plan recommendation systems. We propose a textual representation of our SQL queries and use it to build a *Features Extractor module*. Then, we present a straightforward solution to identify query similarity. This solution relies only on the comparison of the SQL statements. Next, we show how to build an improved solution enabled by machine learning techniques. The improved version takes into account the features of the queries' execution plans. By comparing machine learning algorithms, we find that the improved solution using Classification Based on Associative Rules (CAR) identifies similarity.

Keywords: Recommender System, Machine Learning, Classification, SQL.

I. INTRODUCTION

In this paper, our main objective is to present a recommendation process that uses similarity identification between SQL queries and machine learning techniques to recommend a previously generated execution plan to the optimizer. We firstly present a straightforward approach that relies only on SQL query statements^[1], and presents an improved approach enabled by machine learning techniques afterwards.

In the improved approach, as we represent our queries in two different spaces, we want to use classification methods to try to discover relationships between these representations. Features of the first representation are derived from the SQL statement of the query, while features of the second representation are extracted from the execution plan. More specifically, we are interested in discovering association rules between the textual features of SQL queries and their execution plans' features. These rules are then used to learn models that recommend execution plans for new queries as shown below.




Figure 1. Association rules to plan recommendation.

The remainder of this paper is organized as follows: In the second section, we give an overview of work related to our study. Next, we describe our approach in the third section, and the features extraction process in the fourth section. The straightforward solution and the solution improved by machine learning techniques are presented in Sections 5 and 6, respectively. Finally, we discuss our experimental results in the last section and provide a brief conclusion afterwards.

II. LITERATURE SURVEY

Related work falls into three categories: Machine learning, Recommender systems and Query Optimization. Machine learning techniques are a good alternative to building complex cost-based models and are often easier to adapt and use with new configurations and query types. Accordingly, machine learning methods have been used to tackle a number of issues related to query optimization [3] in general, and SQL queries specifically. These issues range from statistics estimation [4], selectivity and query size estimation [5,6], execution time prediction [7], query performance prediction [8] and execution plan recycling [9,10].

Authors of [11] provide a comprehensive survey on recommender systems and classify them into three categories: 1–content-based, 2–collaborative and 3–hybrid approaches. The concept of recommendation has also been used in the field of query optimization. The work in [12] presents a number of existing methods in the literature to recommend queries to assist the user in the exploration of databases and data warehouses. While the existing approaches using recommendation for query optimization focus on query recommendation; in our work, we are interested in execution plan recommendation.

Existing approaches in the literature perform query similarity identification based on two main phases: (i) defining a pattern for an appropriate query representation; as well as (ii) developing a similarity function. Queries can be represented at the intentional level by considering the uninterpreted SQL sentence [13] or at the extensional level by using the set of tuples resulting from the query execution [14]. Other query representations range from vectors of features [10] to a set of fragments [15]. Graphs [16] are sometimes used as a pattern for query representation as well. The similarity function varies depending on the nature of the problem. [16] uses a simple equality test of query patterns while the comparison in [17] is based on separate

tests of query fragments. Other ways for establishing similarity are based on classical functions applied to the queries representation. For instance, authors of [14] use the inner product, while similarity by [14,18] is identified based on cosine distance and Jaccard similarity, respectively. Our approach for similarity identification is different as we are building two features' datasets, not only one. The first dataset represents the textual features of the queries while the second dataset contains features extracted from their execution plan. Our goal is to find association rules between these two datasets and use them for plan recommendation.

III. METHODOLOGY

In this paper, we propose a mechanism to build an accurate execution plan recommendation system. Figure recommendation plan illustrates our approach.

For a given query q_n , we want the optimizer to choose a suitable execution plan, already stored in the cache memory of the DataBase Management System (DBMS), and reuse it to execute q_n . For this task, we need to study the similarity of q_n with "old" queries that were previously executed by the optimizer. If a similarity is detected between q_n and an "old" query q_o , then the execution plan $Plan_o$ used by the optimizer to execute q_o will be reused to process q_n .

Our study relies on two main components: the first is the "Features Extractor", the task of which is to capture features from query statements and execution plans. The second component is the "Similarity Evaluation" module that receives features as input and predicts similarity. The contributions of this paper can be summarized as follows:

- We firstly propose a textual representation of our data and use it to build the *Features Extractor module*.
- Secondly, we present a straightforward *Similarity Evaluation module* that relies only on the Comparison of the SQL statements of the queries and studies the pitfalls associated with this method.
- Next, we show how to build an improved *Similarity Evaluation module* enabled by machine learning techniques. The improved version of the module takes into account the features of the queries' execution plans.
- We compare three machine learning algorithms and study which one is more accurate in our context.

Feature Extraction

In order to build our *Features Extractor module*, we need to identify a representation model for our data. In the straight forward solution, queries representation (Queries Features Matrix) is derived from their SQL statements, while in the improved solution, an additional representation in the execution plan space is used (Execution Plan Features Matrix). Both matrices are presented in this section.

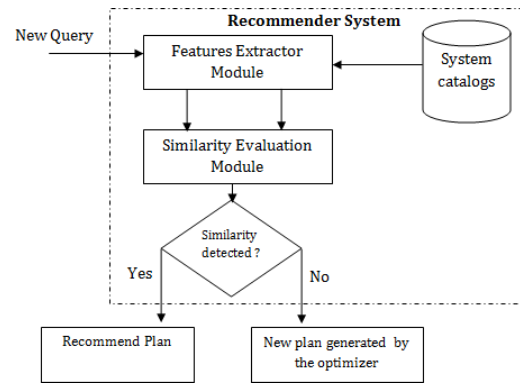


Figure 2. Plan recommendation.

1.1. Query Features Matrix

1.1.1. Features

As a crucial step in processing textual data, text representation converts the content of a given piece of text into a compact format so that it can be recognized and automatically classified. Consequently, the data (e.g., SQL query) is represented as a vector in the feature space, that is, a query Q would be represented as $Q = (F_1, \dots, F_k)$, where k is the feature set size. In our context, a features matrix is an $n \times k$ matrix, where n is the number of query vectors.

Features, more often referred to as *terms*, can be at various levels, such as words, phrases, or any other syntactic indexing units used to identify a SQL query. For our application and approach requirements, we propose a representation method that we call the *Fragments Model*. The fragments model, proposed here, has two main advantages over the classical term model. Firstly, it considerably reduces the dimensions of the feature space and therefore reduces resource consumption in the treatment of queries. This advantage is of crucial importance when we have to deal with query datasets of a large size. The second advantage is the capacity of the fragment model to offer an entity recognition capability. That is, while the classical model treats all tokens of a given query as equivalent terms with the same nature, the fragments model offers the possibility to identify entities in the query. Therefore, relations, attributes and predicates can be recognized that are decisive in identifying queries having similar execution plans.

Fragments Model

In this representation, we split the SQL queries into fragments rather than words. Each fragment is identified by a keyword at the beginning and another one to mark its end in the SQL statement. We have identified six types of fragments as shown in Table 1. Combined fragments of all queries represent the features set.

Table 1. Fragments’ models.

Fragment	Beginning	End
Fragment 1	Select	From
Fragment 2	From	Where, Group By, Order By
Fragment 3	Where	like, =, Group By, Order By
Fragment 4	=	AND, End of the query
Fragment 5	AND	=, End of the query
Fragment 6	like, Group By, Order By	End of the query

Example 1. The fragments’ models of queries Q_0 and Q_1 are presented in Table 2.

Q_0 : SELECT LINE_LEVEL FROM PRODLEVEL WHERE LINE_LEVEL = 'RLI9AYC8YAQ4'

Q_1 : SELECT ACTVARS.DOLLARSALES FROM ACTVARS,PRODLEVEL WHERE ACTVARS.PRODUCT_LEVEL= PRODLEVEL.CODE_LEVEL AND PRODLEVEL.LINE_LEVEL = 'RLI9AYC8YAQ4'

Table 2. Fragments’ models of queries Q_0 and Q_1 .

Fragments	Q_0 Model	Q_1 Model
Fragment 1	LINE_LEVEL	ACTVARS.DOLLARSALES
Fragment 2	PRODLEVEL	ACTVARS,PRODLEVEL
Fragment 3	LINE_LEVEL	ACTVARS.PRODUCT_LEVEL
Fragment 4	'RLI9AYC8YAQ4'	PRODLEVEL.CODE_LEVEL
Fragment 5	NULL	PRODLEVEL.LINE_LEVEL
Fragment 6	NULL	'RLI9AYC8YAQ4'

1.1.2. Weighting Method

Feature weighting is another important step that assigns appropriate weights to features. The simplest weighting method is the binary representation that assigns 1 to a feature if it occurs in the data and 0 otherwise. In this solution, we prefer to use the ‘Term Frequency-Inverse Document Frequency’ (TF-IDF) method rather than the simple frequency weighting. The TF-IDF technique provides a composite weight that is more discriminative compared to a simple term frequency weighting method, and therefore captures similarities in a more precise way. While simple frequency only provides information about presence or absence of a fragment in a given query, TF-IDF gives information on the ‘importance’ of that fragment in the set of queries. That is, the TF-IDF scheme assigns the highest weights to fragments that occur many times within a small number of queries and lowest weights to those occurring in the majority of queries. This property helps in obtaining quality discriminative features, which are not supported by classical term frequency representation.

TF-IDF

TF-IDF is an abbreviation for ‘Term Frequency-Inverse

Document Frequency’. It is a statistical metric that is widely used in information retrieval field to measure the importance of a word in a given document. This measure is calculated using Equations (1) and (2).

Where

$$tf.idf(f, q, Q) = tf(q, f).idf(f, Q),$$

$$idf(f, Q) = \log \frac{N}{|q \in Q : f \in q|}$$

- N : Total Number of Queries in the Database;
- $tf(q, f)$: Frequency of feature f in query q ;
- $|q \in Q : f \in q|$: Number of queries where the feature f appears.

1.2. Execution Plan Features Matrix

In order to build the execution plan features matrix, we need to choose some features from the execution plan of the SQL queries that were previously executed by the optimizer.

The *EXPLAIN PLAN* instruction in *Oracle* stores execution plans in a table called *PLAN_TABLE*. This table contains the number of columns, among which we identify five parameters, namely: Cost, Plan Hash Value, Cardinality, Operation and Bytes. A brief description of the features is given in Table 3, as we use them to represent the SQL query in the execution plan space.

Table 3. Description of the selected execution plan features.

Feature	Description
PLAN_HASH_VALUE	A metric that identifies execution plans.
COST	Execution cost as estimated by the optimizer. The value of this attribute has no measurement unit.
CARDINALITY	Number of rows returned by a query.
OPERATION	Defines the operation type (Select, Update, ...)
BYTES	Number of needed bytes per query.

Straightforward Solution

In this solution, the *Similarity Evaluation module* receives as input a $1 \times k$ features vector V_n of a new query q_n , and an $n \times k$ ‘old’ queries features matrix M . Recall that, n is the number of old queries and k is the number of features. Then, it computes the similarity between V_n and each row of

M . This step produces an $n \times 1$ similarity vector S . Each value $s_{0 \leq i \leq n}$ of S represents the similarity of query $q_{0 \leq i \leq n}$ with q_n . The plan of the query with the highest measure s is then recommended.

Several similarity measures are discussed in the literature. In this paper, we use the cosine similarity, which calculates the cosine of the angle between two feature

vectors as defined in Equation (3), where x_{1k} and x_{2k} are values of feature k in query vectors x_1 and x_2 , respectively.

$$\cos(\theta) = \frac{\sum_{k=1}^n x_{1k}x_{2k}}{\sqrt{\sum_{k=1}^n x_{1k}^2 \sum_{k=1}^n x_{2k}^2}}$$

As we show later in our experiments, this approach proved accurate to recognize queries that are almost identical, but it presents a significant risk to:

- Judge two queries as not similar, by analyzing their SQL syntax, while their execution plans are actually similar.
- Find that two queries have high similarity measure by comparing their statements, but still their plan are different. This is basically due to the fact that small differences in the SQL statements, that could be ignored in classical text categorization tasks, are very decisive in the process of defining an execution plan.

Solution Enabled by Machine Learning Techniques

1.3. Evaluating Classification Techniques

Classification is a basic task in data analysis that requires the construction of a classifier, that is, a function that is “trained” to assign a class label to unknown instances described by a set of features.

In Figure 3, we present the solution enabled by classification techniques. This approach begins with the class assigner module, the mission of which is to evaluate execution plan features of “old” queries and assign same class labels to queries that have equivalent execution plan features. This module’s output is a $1 \times n$ vector C , where each value $c_{0 \leq i \leq n}$ is the class label of query $q_{0 \leq i \leq n}$. The next step is building the training set, which consists of merging vector C to the query features matrix M . Recall that matrix M contains only features that are derived from the SQL statement. Finally, a classifier will “learn” from this training set how to identify similar queries. In this paper, we evaluate three classifiers, namely, Naive Bayes, Support Vectors Machines (SVM) and Classification based on Associative Rules (CAR).

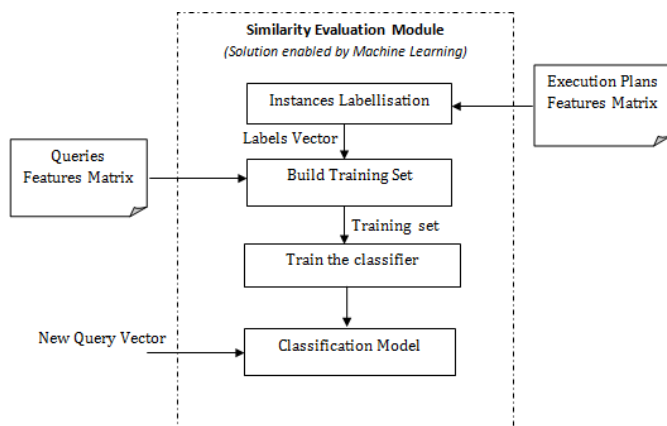


Figure 3. Improved solution enabled by classification

techniques.

1.3.1. Naive Bayes

Naive Bayes classifier builds a model by learning, from a training set, the conditional probability of each attribute x_i given the class label y . Next, classification is performed by applying Bayes rule to compute the probability of y given the particular instance of x_1, \dots, x_n , and then predicting the class with the highest posterior probability. The computation of the predictive model is based on a strong independence assumption that considers all the attributes x_i are conditionally independent given the value of the class y . We evaluate the effect of the Naive Bayes classifier on our data, as it is one of the most effective classifiers with a high predictive performance.

1.3.2. Support Vector Machines (SVM)

Support Vector Machines is basically a two-class classifier that estimates a decision rule $f : R^n \rightarrow \{\pm 1\}$ using training data $(x_1, y_1), \dots, (x_n, y_n) \in R^n \times \{\pm 1\}$ such that f will correctly classify new examples (x, y) . In the case of nonlinearly separable data, SVMs use a kernel function to make the data linearly separable. We could define a kernel $K(x, y)$ as a real-valued function $K : X \times X \rightarrow R$ for which there exists a function $\phi : X \rightarrow F$ with the property:

$$K(x, y) = (\phi(x), \phi(y)).$$

Generally, kernel functions must be positive semi-definite (PSD) to satisfy the theorem of Mercer. Radial Basis Function (RBF) kernel is one of the most efficient and popular kernel functions cited in the literature, thus we use it in our experimentations.

1.4. Evaluating Classification Based on Association Rules

Association rules which were introduced by represent an automatic approach to discover relationships or possible correlations between objects. They were originally designed for finding multi-correlated items in transactions; however, they can be easily adapted for classification.

Thabtah, F gives a comprehensive review of associative classification methods. The main objective of CAR is to find a set of rules that satisfy some minimum support and minimum confidence constraints, and form an accurate classifier. Specifically, in this study, we are interested in finding a set of association rules between two SQL query representations and use them to build an efficient classifier. Recent studies have shown that using classifiers based on class association rules resulted in higher accuracy than those obtained by using other classification algorithms such as C4.5 and ILA, which strengthens our motivation to evaluate CAR in our improved solution.

IV. ANALYSIS & RESULT DISCUSSION

1.1. Experimental Set-Up

1.1.1. Data

Our experimentations were conducted on the Analytical Processing Benchmark (APB-1), that we used to build 2 different datasets:

- Dataset D_1 : Contains 1000 simple queries having different execution plans;
- Dataset D_2 : Contains 1000 queries containing subqueries.

The benchmark contains 55 query templates that we used to build D_1 while we formulated 10 templates to build D_2 .

1.1.2. Models' Validation

We use k -fold cross validation method for model validation. The dataset is divided into k subsets, and the holdout method is repeated k times. Alternatively, each one of the k subsets is used as a test set and the other $k - 1$ subsets are put together to form a training set. Then, the average error across all k trials is computed, and the variance of the resulting estimate is reduced as k increased. We need to consider two aspects when defining the value of parameter k . A large value of k means less bias towards overestimating the true expected error of the model but implies higher variance and higher running time. Hence, we need to choose a value that minimizes the bias without penalizing running time. Generally, the value of k is defined empirically. In our experiments, $k = 10$ offers reasonable bias/variance tradeoff.

1.1.3. Evaluation Metrics

In this paper, we use three metrics in order to measure the prediction accuracy for machine learning techniques.

1-Prediction rate: The rate of correct predictions on the test set. Prediction rate represents the percentage of instances for which predicted class value equals the actual class value.

2-RMSE: Root Mean Square Error (RMSE) is one of the most widely used statistics that measures the differences between values predicted by a model and the actual values.

3-Kappa statistic: Kappa is used as a measure to assess model quality. It is always less than or equal to 1. A value of 1 implies perfect model and values less than 1 imply less than perfect models (0.01–0.2: Slight, 0.21–0.4: Fair, 0.41–0.60: Moderate, 0.61–0.80: Substantial 0.81–0.99: Almost perfect).

1.1.4. Tools

For preprocessing the data, generating and testing the models, we use WEKA, a software developed by the Machine Learning Group at the University of Waikato. We performed the SVM tests with the Sequential Minimal Optimization (SMO) algorithm, which is well-known for its simplicity. We use the Java implementation

(JCBA) of the Classification Based on Associations algorithm (CBA) to evaluate the performance of CAR. In the straightforward solution, Matlab was used to calculate similarity.

1.2. Experimental Results

1.2.1. Experiment 1: Straightforward Solution

In this experiment, we are interested in evaluating the performance of textual models for detecting similarity between SQL queries. We evaluate this solution at two levels:

- Firstly, we compare our fragments model to a word model (with simple frequency weighting) as we want to identify which one of these two models is more likely to identify queries having similar execution plans;
- Secondly, we test the models with our two datasets: D_1 and D_2 .

For each dataset, we choose an arbitrary query to evaluate the solution. Accordingly, for D_1 we consider query Q_{new1} while we take Q_{new2} and Q_{new3} from D_2 . Here, our goal is to see what are the queries that will be identified as similar to Q_{new1} , Q_{new2} and Q_{new3} by this method, using the fragment model. We consider queries for which similarity is higher than 0.95 as similar.

Statements of test queries are as follows:

```
 $Q_{new1}$ : SELECT ACTVARS.DOLLARSALES FROM
ACTVARS.PRODLEVEL WHERE
ACTVARS.PRODUCT_LEVEL=
PRODLEVEL.CODE_LEVEL AND
PRODLEVEL.LINE_LEVEL='RLI9AYC8YAQ4'
```

```
 $Q_{new2}$ : SELECT
CHANNEL_LEVEL,TIME_LEVEL,DOLLARSALES FROM
ACTVARS.CHANLEVEL WHERE
ACTVARS.CHANNEL_LEVEL=
```

```
(SELECT BASE_LEVEL FROM CHANLEVEL WHERE
BASE_LEVEL='C7W3GE3KN9GU')
```

```
 $Q_{new3}$ : SELECT * FROM ACTVARS WHERE TIME_LEVEL
IN (SELECT MONTH_LEVEL FROM TIMELEVEL WHERE
MONTH_LEVEL='199607')
```

Table 5 presents the top three queries (*i.e.*, with the highest similarity values) associated to each test query as returned by the straightforward method. For the purpose of comparison, we additionally calculate, for each of these result queries, the similarities using word model representation. Table 6 reports the query features in the execution plan space.

```
 $Q_{Recom1}$ : SELECT ACTVARS.DOLLARSALES FROM
ACTVARS.PRODLEVEL WHERE ACTVARS.PRODUCT_LEVEL=
PRODLEVEL.CODE_LEVEL
```

```
AND PRODLEVEL.LINE_LEVEL='ITQS4CE29HQR'
```

```
 $Q_{Recom2}$ : SELECT
CHANNEL_LEVEL,TIME_LEVEL,DOLLARSALES FROM
ACTVARS.CHANLEVEL WHERE ACTVARS.CHANNEL_LEVEL=
```

```
(SELECT BASE_LEVEL FROM CHANLEVEL WHERE BASE_LEVEL = 'SMYQ2W6NF5X8')
```

```
Q24: SELECT CHANNEL_LEVEL, TIME_LEVEL, DOLLARSALES FROM ACTVARS, CHANLEVEL WHERE CHANLEVEL.BASE_LEVEL =
```

```
(SELECT CHANNEL_LEVEL FROM ACTVARS WHERE CHANNEL_LEVEL = 'SMYQ2W6NF5X8')
```

```
QRecom3: SELECT * FROM ACTVARS WHERE TIME_LEVEL IN (SELECT MONTH_LEVEL FROM TIMELEVEL WHERE MONTH_LEVEL = '199511')
```

```
Q3Sim: SELECT * FROM ACTVARS WHERE ACTVARS.TIME_LEVEL IN (SELECT TIMELEVEL.MONTH_LEVEL FROM TIMELEVEL
```

```
WHERE TIMELEVEL.MONTH_LEVEL = '199607')
```

Table 5. Cosine similarity/query model.

Dataset	New Query	Suggested Similar Query	Fragments Model	Words Model
D ₁	Q _{new1}	Q _{Recom1}	0.98	0.93
		Q ₁₂	0.90	0.90
		Q ₁₃	0.90	0.89
D ₂	Q _{new2}	Q _{Recom2}	0.97	0.92
		Q ₂₂	0.91	0.90
		Q ₂₃	0.87	0.92
		Q ₂₄	0.33	1
D ₂	Q _{new3}	Q _{Recom3}	0.98	0.93
		Q ₃₂	0.85	0.95
		Q ₃₃	0.78	0.92

Queries identified as similar to Q_{new1} , Q_{new2} and Q_{new3} , using the fragments model, are Q_{Recom1} , Q_{Recom2} and Q_{Recom3} , respectively. By comparing features of queries in the execution plan space (Table 6), we observe that Q_{new1} and Q_{Recom1} are identical, in addition to Q_{new2} and Q_{Recom2} . On the contrary, Q_{Recom3} does not seem to be similar to Q_{new3} as its features are different in the space of execution plans. Moreover, we manually identified and recognized Q_{3Sim} as the actual similar query to Q_{new3} .

On the other hand, we can see that the fragment model is more “discriminant” compared to the word model. Actually, similarity values returned by the word model, for recommended queries, are equivalent to those of the fragments model. However, in the case of Q_{new2} , the word model would have recommended query Q_{24} (instead of query Q_{Recom2}) which is a suboptimal recommendation. In contrast, the fragments model is correctly identifying query Q_{24} as a non similar query with a similarity value that equals 0.33. This result is due to the fact that the fragments model is taking into account the position of the terms in the SQL statements, while the word model is only considering the frequency.

Besides the empirical evaluation, we randomly selected a sample of 100 queries from all datasets and counted the number of queries for which similar queries were identified correctly. Similarity was accurately identified only for the 30% of queries when the word model was applied, while we obtained a rate of 40% with the fragments model. According to the observations above,

while our proposed model is more accurate compared to the simple word frequency model, it still represents a risk to recommend non similar queries.

Table 6. Comparing queries in the execution plan space.

Query	Cardinality	Bytes	Plan Hash Value
Q _{new1}	196,714	9,638,986	1,918,545,856
Q _{Recom1}	196,714	9,638,986	1,918,545,856
Q ₁₂	688,500	42,687,000	1,918,545,856
Q ₁₃	688,500	42,687,000	1,918,545,856
Q _{new2}	81,000	2,187,000	4,072,885,321
Q _{Recom2}	81,000	2,187,000	4,072,885,321
Q ₂₂	81,000	7,695,000	4,072,885,321
Q ₂₃	79,442	1,906,608	4,072,885,321
Q ₂₄	65,548	1,713,152	4,072,885,321
Q _{new3}	4673	387,859	2,685,087,043
Q _{Recom3}	4490	372,670	2,685,087,043
Q _{3Sim}	4490	372,670	2,685,087,043
Q ₃₂	68,850	3,167,100	2,685,087,043
Q ₃₃	68,850	3,167,100	2,685,087,043

Experiment 2: Solution Enabled by Machine Learning Techniques.

The second experiment was conducted on D_2 dataset.

Table 7 reports the prediction rates of SVM, Naive Bayes and CAR algorithms. We observe that the highest rate is realized by the JCBA algorithm followed by Naive Bayes and SVM.

Table 7. Comparison of prediction rate.

SVM	NB	CAR
77.45%	78.65%	91.5%

Figure 4 provides a comparison of the Kappa statistic and the RMSE of the classifiers. We can see that the minimal value of RMSE (0.28) is given by JCBA while Naive Bayes and SVM perform almost in the same way with values equal to 0.48 and 0.41, respectively. Furthermore, evaluation of the Kappa statistic reveals that JCBA algorithm is generating an “Excellent” model while Kappa values of Naive Bayes and SVM are in the same range and provide a “Fair” Model.




Figure 4. Comparison of Kappa and RMSE.

While the straight forward solution was not able to provide an accurate results for dataset D_2 , the improved solution using machine learning techniques performed much better with a prediction rate of 76.55%. Consequently, better recommendation results are obtained, particularly when using Classification based on Association Rules.

V. CONCLUSIONS

In this paper, we presented an approach for execution plan recommendation based on similarity identification. Our goal was to build a model that recognizes similarity by using only the SQL statement of new queries. We used machine learning techniques to improve the similarity detection. Classification based on Association Rules has better potential to provide accurate prediction compared to SVM and Naive Bayes. Moreover, the proposed model was able to identify and associate similar queries having similar execution plans. Therefore, our approach is able to recommend relevant execution plans to new queries, based on similarity identification with historical queries that were previously executed by the optimizer. While execution time of models could be a constraint for using CAR, working with a parallelized version of the JCBA implementation, in future studies, can help resolve this issue. Another perspective is to use adaptive recommendations using the optimizers' feedback to enhance the prediction accuracy.

VI. REFERENCES

- [1] Zahir J.; Qadi A.; Recommendation System for Execution Plans Using Machine Learning. *Math. Comput. Appl.* 2016, 21, 23; doi:10.3390/mca21020023.
- [2] Ghosh, A.; Parikh, J.; Sengar, V.S.; Haritsa, J.R. Plan selection based on query clustering. In *Proceedings of the 28th International Conference on Very Large Data Bases*, Hong Kong, China, 20–23 August 2002.
- [3] Adomavicius, G.; Tuzhilin, A. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. Knowl. Data Eng.* 2005, 17, 734–749.
- [4] Marcel, P.; Negre, E. A Survey of Query Recommendation Techniques for Data Warehouse Exploration; *Entrepôts de Données et l'Analyse en ligne—EDA*: Clermont-Ferrand, France, 2011; pp. 119–134.
- [5] Yao, Q.; An, A.; Huang, X. Finding and analyzing database user sessions. In *Database Systems for Advanced Applications*; Springer: Berlin/Heidelberg, Germany, 2005; pp. 851–862.
- [6] Stefanidis, K.; Drosou, M.; Pitoura, E. You May Also Like results in relational databases. In *Proceedings of the International Workshop on Personalized Access, Profile Management and Context Awareness: Databases*, Lyon, France, 24–28 August 2009; Volume 9.
- [7] Wu, W.; Chi, Y.; Hacigümüs, H.; Naughton, J.F. Towards predicting query execution time for concurrent and dynamic database workloads. *Proc. VLDB Endow.* 2013, 6, 925–936.
- [8] Selinger, P.G.; Astrahan, M.M.; Chamberlin, D.D.; Lorie, R.A.; Price, T.G. Access path selection in a relational database management system. In *Proceedings of the 1979 ACM SIGMOD International Conference on the Management of Data*, Boston, MA, USA, 30 May–1 June 1979.
- [9] Hasan, R.; Gandon, F. A machine learning approach to SPARQL query performance prediction. In *Proceedings of the 2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, Warsaw, Poland, 11–14 August 2014; IEEE: Piscataway, NJ, USA, 2014; Volume 1, pp. 266–273.
- [10] Gao, L.; Wang, M.; Wang, X.S.; Padmanabhan, S. A learning-based approach to estimate statistics of operators in continuous queries: A case study. In *Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, San Diego, CA, USA, 13 June 2003; ACM: New York, NY, USA, 2003; pp. 66–72.
- [11] Gryz, J.; Liang, D. Query selectivity estimation via data mining. In *Intelligent Information Processing and Web Mining*; Springer: Berlin/Heidelberg, Germany, 2004; pp. 29–38.
- [12] Chaudhuri, S.; Ganti, V.; Gravano, L. Selectivity estimation for string predicates: Overcoming the underestimation problem. In *Proceedings of the 20th International Conference on Data Engineering*, Boston, MA, USA, 30 March–2 April 2004; IEEE: Piscataway, NJ, USA, 2004; pp. 227–238.
- [13] Ganapathi, A.; Kuno, H.; Dayal, U.; Wiener, J.L.; Fox, A.; Jordan, M.; Patterson, D. Predicting multiple metrics for queries: Better decisions enabled by machine learning. In *Proceedings of the IEEE 25th International Conference on Data Engineering, ICDE'09*, Shanghai, China, 29 March–2 April 2009; IEEE: Piscataway, NJ, USA, 2009; pp. 592–603.
- [14] Akdere, M.; Çetintemel, U.; Riondato, M.; Upfal, E.; Zdonik, S.B. Learning-based query performance modeling and prediction. In *Proceedings of the 2012 IEEE 28th International Conference on Data Engineering (ICDE)*, Washington, DC, USA, 1–5 April 2012; IEEE: Piscataway, NJ, USA, 2012; pp. 390–401.
- [15] Sarda, P.; Haritsa, J.R. Green query optimization: Taming query optimization overheads through plan recycling. In *Proceedings of the Thirtieth International Conference on Very Large Data Bases*, Toronto, ON, Canada, 31 August–3 September 2004.
- [16] Aligon, J.; Golfarelli, M.; Marcel, P.; Rizzi, S.; Turricchia, E. Mining preferences from OLAP query logs for proactive personalization. In *Advances in Databases and Information Systems*; Springer: Berlin/Heidelberg, Germany, 1998; pp. 84–79.
- [17] Yang, X.; Procopiuc, C.M.; Srivastava, D. Recommending join queries via query log analysis. In *Proceedings of the IEEE 25th International Conference on Data Engineering, ICDE'09*, Shanghai, China, 29 March–2 April 2009; IEEE: Piscataway, NJ, USA, 2009; pp. 964–975.
- [18] Gupta, A.; Mumick, I.S. *Materialized Views: Techniques, Implementations, and Applications*; MIT Press: Cambridge, MA, USA, 1999.
- [19] Chatzopoulou, G.; Eirinaki, M.; Koshy, S.; Mittal, S.; Polyzotis, N.; Varman, J.S.V. The QueRIE system for Personalized Query Recommendations. *IEEE Data Eng. Bull.* 2011, 34, 55–60.