

Hanyin Cheng

East China Normal University, China

hanyinch@gmail.com

Bin Yang

East China Normal University, China

byang@dase.ecnu.edu.cn

Christian S. Jensen

Aalborg University, Denmark

csj@cs.aau.dk

RCRank: Multimodal Ranking of Root Causes of Slow Queries in Cloud Database Systems

Biao Ouyang East China Normal University, China bouyang@stu.ecnu.edu.cn

Yang Shu* East China Normal University, China yshu@dase.ecnu.edu.cn

> Qingsong Wen Squirrel Ai Learning, USA qingsongedu@gmail.com

Yingying Zhang Alibaba Cloud Computing, China congrong.zyy@alibaba-inc.com

Chenjuan Guo East China Normal University, China cjguo@dase.ecnu.edu.cn

Lunting Fan Alibaba Cloud Computing, China lunting,fan@taobao.com

ABSTRACT

With the continued migration of storage to cloud database systems, the impact of slow queries in such systems on services and user experience is increasing. Root-cause diagnosis plays an indispensable role in facilitating slow-query detection and revision. This paper proposes a method capable of both identifying possible root cause types for slow queries and ranking these according to their potential for accelerating slow queries. This enables prioritizing root causes with the highest impact, in turn improving slow-query revision effectiveness. To enable more accurate and detailed diagnoses, we propose the multimodal Ranking for the Root Causes of slow queries (RCRank) framework, which formulates root cause analysis as a multimodal machine learning problem and leverages multimodal information from query statements, execution plans, execution logs, and key performance indicators. To obtain expressive embeddings from its heterogeneous multimodal input, RCRank integrates self-supervised pre-training that enhances cross-modal alignment and task relevance. Next, the framework integrates rootcause-adaptive cross Transformers that enable adaptive fusion of multimodal features with varying characteristics. Finally, the framework offers a unified model that features an impact-aware training objective for identifying and ranking root causes. We report on experiments on real and synthetic datasets, finding that RCRank is capable of consistently outperforming the state-of-the-art methods at root cause identification and ranking according to a range of metrics.

PVLDB Reference Format:

Biao Ouyang, Yingying Zhang, Hanyin Cheng, Yang Shu, Chenjuan Guo, Bin Yang, Qingsong Wen, Lunting Fan, and Christian S. Jensen. RCRank: Multimodal Ranking of Root Causes of Slow Queries

in Cloud Database Systems. PVLDB, 18(4): 1169 - 1182, 2024. doi:10.14778/3717755.3717774

Proceedings of the VLDB Endowment, Vol. 18, No. 4 ISSN 2150-8097. doi:10.14778/3717755.3717774

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at https://github.com/decisionintelligence/RCRank.

1 INTRODUCTION

Companies and individuals are increasingly migrating their database services to the cloud. At the same time, poor cloud database system performance in the form of slow queries causes economic losses to users and decreases user trust in cloud-based data management. For example, every 0.1-second increase in Amazon's page load time causes a 1% drop in sales, and a 0.5-second increase in Google's search latency causes a 20% drop in traffic [18]. Thus, accelerating slow queries is essential for ensuring high-performance cloud database systems.

Slow queries may be attributed to database system internal factors, such as the absence of relevant indexes or poorly written SQL statements, and external factors, such as I/O bottlenecks and network issues. In this paper, we aim at providing a framework to enable users to address slow queries while focusing on root causes deriving from internal factors. This is because external factors are often beyond the control of users and require the intervention of cloud database service providers. In contrast, internal factors can be controlled by cloud database users, such as by modifying indexes or rewriting queries.

Identifying root causes, i.e., factors that cause slow queries, allows for performance improvement through corresponding revision methods. Thus, root cause analysis identifies the causes of slowquery types through root cause classification [23, 51, 52]. While methods exist that target slow-query identification, two limitations remain:

Identification of Root Causes Mainly. Existing methods [23, 44, 52] focus on identifying root causes, as shown in Figure 1. This does not fully support the need to focus on addressing the "most significant" root causes. Revising slow queries based on root causes is costly, and making revisions according to every root cause can result in substantial expenses. Therefore, it is attractive to take into account the impact of root causes (i.e., how much runtime can be saved if a root cause is addressed) when choosing which one to address. Root Cause Identification (RCI) fails to quantify the

^{*}Corresponding author

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit https://creativecommons.org/licenses/by-nc-nd/4.0/ to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

potential impact of addressing identified root causes, thus limiting the practical utility.

Incomplete Observability of Cloud Database Systems. As shown in Figure 1, most methods [23, 44] rely on single-modal information, e.g., key performance indicators, such as CPU or memory use time series, to identify root causes, disregarding other data sources that can offer insight into internal factors related to slow query processing. For example, query statements and execution plans contain information about query objectives and estimated execution processes; and execution logs contain information about resources consumed during query execution and execution status. Aiming for full observability by taking all these data sources into account offers a better foundation for understanding slow queries and their causes, thereby enabling more accurate root cause identification.

To eliminate the two limitations, we propose ranking the impact of the root causes using multimodal observability data.

Ranking the impact of root causes. We propose a model to estimate the impact of root causes, thus providing a ranked list of identified root causes. To unify identification and ranking, we propose a training objective for ranking root causes by estimating their impact. Two impact-aware regularizations are utilized in the training objective to enhance the identification of valid root causes and preserve the order of their impacts. Specifically, we collect slow queries from the cloud database system that encompass query statements, execution plans, execution logs, and key performance indicators. We then revise the queries according to potential root causes and record the relative improvement of the revised queries as the impact of each root cause. For instance, the runtime of a given slow query may drop by 14.8%, after updating Outdated Statistics. Then we associate 14.8% with the "Outdated Statistics" root cause for the query. The multimodal information about queries, along with the impact of root causes are used for training a model. During inference, the model (i) diagnoses slow queries to identify root causes that lead to the slow queries and (ii) ranks the root causes according to their potential impact.

Encoding multimodal representations. We encode the available multimodal information into a unified embedding to support root cause ranking. The information encompasses query statements in text form, execution plans represented as graphs, logs in keyvalue format, and key performance indicators (KPIs) as time series. To integrate this information, we first use different encoders to encode each modality into embeddings. Then we employ a fusion embedding module to encode the heterogeneous embeddings of the different modalities into a shared embedding space for further feature extraction. We propose self-supervised pre-training to enhance cross-modal alignment and task relevance to obtain more expressive embeddings. To enable effective fusion of multimodal features, we propose a multimodal fusion model that uses rootcause-adaptive cross Transformers, capable of adaptively fusing multimodal features for the diagnosis of root causes with varying characteristics.

To the best of our knowledge, this is the first multimodal method for ranking the root causes of slow queries for root cause diagnosis. The contributions are summarized as follows.







(b) Root Cause Ranking (RCR)

Figure 1: Root cause identification (RCI) vs. root cause ranking (RCR). (i) RCI often utilizes partial observability, whereas RCR utilizes multimodal, full observability. (ii) RCI only identifies possible root causes, whereas RCR ranks root causes according to their potential impact, enabling users to identify the most significant root causes.

- We propose means to automatically collect root cause impact labeled data, which provides a data foundation for multimodal root cause diagnosis.
- We propose a multimodal representation method that is capable of harnessing full observability by implementing expressions for embedding heterogeneous modalities, modality alignment, and task relevance, while effectively and adaptively fusing multimodal features.
- We propose a root cause impact ranking method that enables unified identification and ranking of root causes.
- We consider datasets from the Alibaba Hologres system as well as synthetic datasets. Experimental results show that our method is capable of outperforming existing approaches in terms of root cause identification and ranking according to a range of metrics.

2 RELATED WORK

2.1 Slow Query Analyses

We classify slow-query analysis methods in Table 1 from two perspectives: root cause identification vs. ranking for slow queries and single-modality vs. multi-modality modeling.

Existing Root Cause Identification (**RCI**) methods [11, 15, 23, 44, 52] use clustering, attribution, or prediction methods to identify root causes. However, these methods cannot estimate or rank the impact of root causes (**RCR**). For example, when diagnosing root causes of slow queries, multiple root causes can contribute to an issue, so

Table 1: Categorization of slow query diagnosis.

	Root Cause Identification (RCI)	Root Cause Ranking (RCR)
Single-Modality Observability	iSQUAD [23] ExplainIt! [11] DBSherlock [44]	-
Multi-Modality Observability	OpenGauss [15] D-BOT [52]	RCRank (ours)

that addressing each can lead to different levels of performance improvement. Users are interested primarily in root causes that can bring the most significant performance gains.

Previous studies primarily focus on slow-query analysis through **single-modal input**, e.g., key performance indicators. DBSherlock [44] focuses on the precise values of KPIs. A specific KPI value can only suggest whether or not an indicator is anomalous. iSQUAD [23] focuses on the root causes of intermittently slow queries and can effectively inform DBAs of the specific category of an anomaly. However, when iSQUAD encounters root causes outside its clustering, manual analysis of the anomaly types of KPIs is required. ExplainIt! [11] employs a probabilistic graphical model facilitating causal reasoning to analyze root causes. Interpretable ML methods [9, 23, 44] focus on root cause analysis through KPIs and offer reasonable interpretations.

Two existing methods fuse **multimodal input** encompassing system metrics and query information to identify root causes. Open-Gauss [15] discovers time-consuming operators in SQL by fusing execution plans and system KPI metrics. D-Bot [52] leverages large language models [29, 39, 43, 53, 54] alongside maintenance knowledge to analyze monitoring and query information, facilitating the identification of root causes of slow queries. However, in our scenario, additional observable information sources are available. Since users prioritize addressing root causes of slow queries with high impact, we not only identify the root causes but also rank these according to their impact to facilitate effective slow-query resolution. Our study considers query statements, execution plans, execution logs, and time series to rank root causes of slow queries according to their impact.

2.2 Multimodal Learning

Previous multimodal learning methods [34, 50] focus on the fusion of text, images, and audio. The fusion typically involves feature fusion and alignment through concatenation and cross-attention. Tokens from multiple modalities are concatenated directly, adding positional information to distinguish the modality of each input. The aim is to understand the context of other modalities. However, this approach comes with high computational complexity. Instead, we use the cross-attention mechanism to reduce computational complexity. Cross-Attention [16, 17, 20, 21, 28] exchanges information between modalities. However, the large number of modalities may still result in an increase in computational complexity when pairwise calculations are involved. We thus use the primary modality to avoid pairwise interactions between all modalities.

When fusing text, images, and audio, existing methods extract commonalities and specificities across the modalities [10, 45, 46]. The methods achieve this by using shared modules for extracting common features across text, images, and audio, while using individual modules to capture modality-specific characteristics. AnoFusion [47] models the failure instance detection problem by capturing the relationship between time series, graphs, and text. Our method focuses on the complex correlations among text, graphs, and time series [3–5, 31, 37, 40, 41, 48] to rank the root causes of slow queries.

Additionally, some methods focus on improving query performance, with deep learning representation methods pre-training on SQL and query plans [36, 49], and encoding plan methods [1, 25, 26, 49]. Some methods encode SQL for cost estimation [12, 35, 36] and query optimization [2, 13, 22, 24, 55], and some method [30] encodes metrics for autoscaling. Some methods focus on enhancing query performance through the automatic adjustment of indexes [6, 8, 14, 32, 33]. In contrast, we utilize more extensive information for identifying and ranking root causes of slow queries.

3 PROBLEM SETTING AND FORMULATION

3.1 Query Formalization

A query $X_i = (S, P, L, I)$ is a four-tuple composed of a query statement *S*, an execution plan *P*, an execution log *L*, and KPIs *I*. We define this as follows.

- Query Statement (SQL): A query statement is the SQL statement of the query, which is a sequence S = (s₁, s₂, ..., s_n), where s_i is an SQL text token (e.g., "Select"). This captures the intent of the query and relevant database information, e.g., table names.
- Execution Plan (Plan): An execution plan is a Directed Acyclic Graph (DAG) P = (V, E), where V is a set of operators and a directed edge e_k from v_i to v_j exists if operator v_j is the next operation of v_i.
- Execution Log (Log): Log files record information related to queries, e.g., the query duration and number of rows read by the execution of the query. In our setting, logs are represented as vectors $L = (l_1, l_2, ..., l_m)$, where l_i is a record related to the query.
- Key Performance Indicators (KPIs): The execution times of query statements are influenced by the current performance state of the database. KPIs $I = (I_1, I_2, ..., I_q)$ are composed of qmulti-dimensional time series, each of length t, that are employed for evaluating database performance. Here q is the number of evaluated metrics. These include memory usage percentages, CPU utilization percentages, and I/O counts.

3.2 Root Causes for Slow Queries

Many factors can lead to slow queries, including external factors such as CPU and memory depletion, resource contention leading to deadlocks, and instance node crashes. Internal factors such as suboptimal database structures and poorly-written query statements also contribute to slow queries. While extensive research exists on diagnosing the root causes of slow queries related to external factors, database users are often interested primarily in the revisions they can make in terms of database structures and query statements to speed up query execution.

We define a **root cause** as an internal factor within the database, e.g., a slow query caused by sub-optimal database structures or poorly-written query statements. For example, the root cause of *missing indexes* means that, in scenarios with a large volume of data, queries without indexes result in full table scans, consuming substantial computational resources and time. Adding appropriate indexes can enhance query efficiency. By employing suitable solutions to address root cause issues, we can improve the execution time of slow queries. After revising slow queries using revision methods corresponding to different root causes, the query performance is improved. Users are particularly concerned with root causes whose revision methods can contribute the most to improving query performance, e.g., the root cause that, when addressed, improve query execution time the most.

The impact of a root cause captures how much a slow query is influenced by the root cause. More precisely, consider a slow query X_i and potential root causes $RC = \{RC_1, RC_2, ..., RC_r\}$, where *r* is the total number of root cause types. We define an impact value y_{ij} that captures how much the performance of the query can be improved when revised according to root cause RC_j :

$$y_{ij} = \frac{runtime(X_i) - runtime(revise(X_i, RC_j))}{runtime(X_i)}$$
(1)

Here, $revise(X_i, RC_j)$ denotes the original slow query X_i revised according to root cause RC_j . As an example, the original runtime of X_i is 1.5 seconds, and RC_j indicates missing indexes. Then $revise(X_i, RC_j)$ involves creating relevant indices for query X_i , and the revised runtime $runtime(revise(X_i, RC_j)$ becomes 0.67 seconds. Therefore, the impact y_{ij} of RC_j on X_i , is (1.5 - 0.67)/1.5 = 55.3%, meaning that after revision w.r.t. root cause RC_j , the runtime of the slow query X_i is improved by 55.3%.

Valid root causes A root cause RC_j is valid for slow query X_i if y_{ij} exceeds a threshold ϵ .

3.3 **Problem Formulation**

Given a root cause impact dataset $D = \{(y_{ij}, X_i, RC_j)\}$ consisting of impacts of root causes for different slow queries, we aim at learning a model f that can estimate the impact \hat{y}_{kj} for an unseen query X_k for each root cause RC_j . This can be formulated as follows.

$$\hat{y}_{kj} = f_{\theta}(X_k, RC_j), \tag{2}$$

where θ represents the learnable parameters for the estimation model *f* and $1 \le j \le r$.

4 MULTIMODAL DIAGNOSIS FRAMEWORK

4.1 Overview

We propose a multimodal diagnosis framework for identifying root causes of slow queries and ranking them by their impact. As shown in Figure 2, the framework consists of two main modules: a slow query and root cause collection module, and a multimodal root cause diagnosis module.

The **slow query and root cause collection module** includes cloud database system monitoring, slow query collection, and root cause collection. The module collects slow queries and the impacts of root causes on the slow queries, which serves as a data foundation for the second module. The **multimodal root cause diagnosis module** learns a mapping from unseen slow queries X_k and root causes RC_j to the estimated impact \hat{y}_{kj} of RC_j on X_k , thus facilitating the construction of lists of root causes according to their estimated impact.



Figure 2: Overview of the multimodal diagnosis framework for root causes of slow queries.

4.2 Slow Query and Root Cause Collection

We collect slow queries and the impact of their corresponding root causes as $\{(y_{ij}, X_i, RC_j)\}$. As shown in Figure 2, slow query collection involves continuous monitoring of queries and database instances through a cloud database monitoring system, collecting queries that exceed the slow query threshold δ . Root cause collection involves obtaining slow query root causes through analysis by a rule-based method and a large language model (LLM) based method.

Slow query collection continuously monitors the execution times of queries through the cloud database monitoring system, collecting the query statements (SQL), execution plans (Plan), and execution logs (Log) of the queries. The performance monitoring system continually gathers instance metric information (KPIs), such as the memory usage percentage and CPU utilization percentage, thereby obtaining the status. We collect all monitored queries { Q_i } for pre-training, with queries exceeding the slow query threshold δ considered as slow queries.

Root cause collection involves analyzing slow queries and revising slow queries according to different root causes. Therefore, we design two types of automated methods: *rule-based* and *LLM-based methods*. For root causes that can be addressed easily by employing predefined SQL templates, we choose rule-based methods. For the remaining cases, we use LLM-based methods.

Rule-based methods. As shown in Figure 2, the rule-based method includes three steps. First, according to a specific root cause RC_j , we employ corresponding rules to revise the slow query using predefined SQL templates. Second, we re-execute the slow query X_i to obtain the execution time after revision. Third, we calculate the impact y_{ij} according to Equation 1. We detail the rule-based methods for two specific root causes for illustration.

 \star Statistics. Unsynchronized statistical information can result in the generation of sub-optimal execution plans, thus causing slow queries. For a slow query, we extract all tables involved in the query and then update the statistics by executing the SQL "ANALYZE tablename" for these tables. After updating the statistics, we reexecute the query and use Equation 1 to compute the impact of the root cause "Statistics."

★ Join-Order. When SQL join relationships are complex or involve many tables, the optimizer spends considerable time selecting the optimal join order. To collect slow queries that require adjusted join strategies, we re-execute the slow queries by modifying the execution strategy. For example, we can utilize "set optimizer_join_order = greedy;" in SQL to adjust the execution strategy. Alternatively, we can use the "exhaustive" or "query" strategies. We record the runtime of the best join order, i.e., the join order with the shortest runtime. Based on this, we then compute the impact of the root cause "Join-Order."

LLM-based methods. As shown in Figure 2, the LLM-based method encompasses four steps. First, we input query information, database information, and root cause types, and then we generate a prompt [19] based on the prompt template shown in Figure 3. Second, based on the generated prompt, the LLM recommends slow-query revisions. Third, we re-execute the revised query to obtain the execution time after revising the query according to root cause RC_j . Fourth, we calculate the impact y_{ij} according to Equation 1. The three root causes addressed using LLM-based methods are handled as follows.

★ Index. With no indexes available, the database stem executes queries by scanning entire tables row by row to find matching rows. On the other hand, too many indexes can result in index inefficiency. We employ the LLM approach to recommend indexes that yield the shortest execution time. First, we provide the original query statement, execution plan, execution logs, and table structure information as a prompt; see Figure 3. Then, the LLM recommends indexes based on the prompt. We compare the execution times when using the indexes recommended by the LLM and those when using the original indexes to determine the impact of the root cause "Index."

★ Distribution-key. The distribution key facilitates the partitioning of data into shards, and a uniform distribution helps prevent data skew. We provide the original query statement, execution plan, execution logs, and table structure information as a prompt; again, see Figure 3. We compare the execution time of the distribution keys recommended by the LLM with those of the original distribution keys to determine the impact of the root cause "Distribution-key."

★ Query. Queries written by users may contain redundant operators, nested subqueries, etc., so that reduced execution times can result from rewriting the queries into equivalent queries. We provide the original query statement, execution plan, execution logs, and table structure information as a prompt; again, see Figure 3. The LLM recommends rewritten queries based on the prompts. Considering that the LLM may output incorrect queries, we inform the LLM about erroneous queries along with the corresponding error messages to get better recommendations. We compare the execution times of rewritten queries with those of the original query to determine the impact of the root cause "Query."

By collecting slow queries and revising them according to methods associated with corresponding root causes, we construct a dataset $D = \{(y_{ij}; X_i, RC_j)\}$ for identifying and ranking root causes of slow queries.

Context: Imagine you're the database administrator for a big-sized e-commerce company that stores customer and product information in Postgres database.	
SQL: {query}	
Execution plan: {execution plan}	
Execution logs: {execution logs}	
Tbale columns: {table columns}	
Index/Distribution-key Revision: You will recommend indexes/(distribution keys) from	
the table columns based on the SQL, execution plan, and execution logs to reduce SQL execution time.	
Query Revision: You will rewrite SQL based on the above SQL, execution plan, execu-	
tion logs, and table structure to reduce SQL execution time. Make sure the SQL can be	

Figure 3: Prompt template with context and revision sections for LLM-based slow-query revision.

4.3 Multimodal Root Cause Diagnosis Model

We propose a multimodal diagnosis model for identifying the root causes of slow queries and ranking the impact of different root causes based on dataset $D = \{(y_{ij}, X_i, RC_j)\}$. A overview of the framework and its three main modules is shown in Figure 4

The input embedding module incorporates an encoder for each input modality. Given the multimodal input data X_i consisting of a query statement S, an execution plan P, an execution log L, and key performance indicators I, the encoders output the embeddings of each modality $\mathbb{E}(X_i) = (E_S, E_P, E_L, E_I)$. The module determines how we encode the heterogeneous inputs of different modalities into a shared embedding space for further feature extraction. Note that the proposed pre-training processes are self-supervised [42], which reduces the amount of root-cause-annotated data needed for supervised training and enables the utilization of unlabeled query data for obtaining more expressive embeddings. In Section 4.3.1, we describe the design of the input embedding module. To improve the quality of embeddings for better cross-modal alignment and task relevance and to reduce the amount of root-cause-annotated data needed for training, we employ self-supervised multimodal pre-training for the input embeddings.

The **multimodal fusion** module takes the encoded embeddings of each modality as input. The module learns to fuse information from these embeddings to extract the features $E_F = f_{\text{fusion}}(E_S, E_P, E_L, E_I)$ for root cause diagnosis. To facilitate feature extraction in the module, we propose a cross-modal Transformer architecture in Section 4.3.2. We design feature decomposition and adaptive gating mechanisms to enable adaptive multimodal fusion for different root causes.

The **root cause estimation** module leverages the extracted multimodal fusion features to obtain the final estimations $Y_{RC} = f_{\text{pred}}(E_F)$ regarding the root causes. This module's training and inference processes are designed to identify and rank the impacts of root causes. It includes a learning objective for estimating the impacts of different root causes to identify and rank root causes in a unified model. The estimation accuracy is enhanced by impactaware regularization for valid and highly impactful root causes. This is covered in Section 4.3.3.

4.3.1 Self-supervised Pre-trained Embeddings.

Input Embeddings. The input embedding module encodes inputs with different modalities. Given the multimodal input data X_i , we employ different encoders tailored for the different modalities in the input (denoted as Enc) to obtain an embedding for each single



Figure 4: Overview of the multimodal learning model for root cause diagnosis, which is composed of three main modules: (1) input embedding module, (2) multimodal fusion module, and (3) root cause estimation module.

modality in a common embedding space:

$$(E_S, E_P, E_L, E_I) = \mathbb{E}_{\phi}(X_i)$$

= (Enc_S(S), Enc_P(P), Enc_L(L), Enc_I(I)). (3)

More specifically, we employ BERT [7] as the encoder of query statements *S*, QueryFormer [49] as the encoder of execution plans *P*, a multi-layer perceptron as the encoder of execution logs *L*, and a 2DCNN model as the encoder of KPIs *I*.

Self-supervised Multimodal Pre-training. In the multimodal data inputs, partial correspondence exists among the different modalities. For example, field names, table names, and filter conditions in query statements also occur in nodes in the execution plan tree. Likewise, some data in execution logs has corresponding data in the root nodes of plans. Inspired by such correspondences in the multimodal query data, we propose self-supervised pre-training to align the different modalities and enhance the expressiveness of the data embeddings.

As shown in Figure 5, to align query statements, execution plans, and execution logs, we propose to mask critical information within each modality and then learn to predict the masked information jointly using the other modalities. We first match the information of query conditions, join conditions, table names, column names, and operations between the query statement and execution plan. Additionally, we match specific numerical information between the execution log and the execution plan. Then, we randomly select some matched parts as the critical information and mask the information in one of the modalities. Specifically with a 50% probability, we choose to mask tokens in either the query statement or the execution plan. The corresponding tokens will be masked and serve as reference prediction information. For the execution logs, we mask parts of the logs and utilize information from query statements and execution plans to predict the masked parts of the logs. This approach enables the model to align query statements, execution plans, and execution logs at the encoding level. For example, we replace the table name "user" in the query statement with "<mask>"



Figure 5: Illustration of self-supervised pre-training for multimodal encoders *Enc*_S, *Enc*_O, *Enc*_L, and *Enc*_I.

while keeping the "user" unchanged in the execution plan. By using the masked query statement, execution plan, and execution logs, we predict the masked token. Following this, we employ a transformer module *A* to aggregate the three modalities and jointly predict the masked content. Take the SQL modality as an example and let \tilde{S} denote its masked content and $S - \tilde{S}$ denote its remaining content. The joint prediction objective for self-supervised pre-training is applied to each modality to achieve fine-grained cross-modal alignment:

$$\mathcal{L}_{SQL} = \left\| \tilde{S} - A \left(\text{Enc}_{S}(S - \tilde{S}), \text{Enc}_{P}(P), \text{Enc}_{L}(L) \right) \right\|_{2}^{2}$$

$$\mathcal{L}_{PLAN} = \left\| \tilde{P} - A \left(\text{Enc}_{S}(S), \text{Enc}_{P}(P - \tilde{P}), \text{Enc}_{L}(L) \right) \right\|_{2}^{2} \qquad (4)$$

$$\mathcal{L}_{LOG} = \left\| \tilde{L} - A \left(\text{Enc}_{S}(S), \text{Enc}_{P}(P), \text{Enc}_{L}(L - \tilde{L}) \right) \right\|_{2}^{2}$$



Figure 6: Architecture of the cross-modal Transformer.

The KPIs modality contains information about the current database state. Considering the impact of the database state on query execution, it is crucial to capture both temporal features of the database state and the patterns of interaction between different metrics of the database instance. KPIs represent the database state, and we explore their temporal characteristics. To capture patterns in metrics, we pre-train the embeddings of metrics separately. Inspired by this property, we propose to employ an auto-encoding structure and reconstruction of the KPI time series as the pre-training objective:

$$\mathcal{L}_{\text{KPIs}} = \|\text{Dec}_I(\text{Enc}_I(I)) - I\|_2^2,$$
(5)

where Dec_I is the decoder that is to reconstruct the original KPI input from its embeddings. Through this pre-training process, we can obtain better embeddings for KPIs related to anomaly detection, which are very relevant to the downstream task of slow queries, i.e., estimating the impact of root causes.

Pre-training with large quantities of queries: As the pretraining of the encoders is self-supervised and only uses the queries themselves, not their root cause impact, we can use all queries, not only the slow queries, in the root cause impact dataset *D*. Thus, we use a large quantity of queries collected from the cloud database system to facilitate the pre-training.

4.3.2 Root-Cause-Adaptive Cross Transformer.

Cross-modal Transformer. To more fully utilize complementary information from the diverse modalities, the multimodal fusion module learns to fuse them to extract features for root cause diagnosis. We employ the Transformer [38] architecture and cross-modal attention to integrate information from multiple modalities.

Cross-attention (*CA*) [17] uses the same computational method as self-attention to process the relationship between different input data A and B:

$$CA(A, B) = \operatorname{softmax}\left(\frac{AW_A^Q BW_B^K}{\sqrt{d_q}}\right) BW_B^V,$$
(6)

where W_A^Q is the query projection layer of A, W_B^K is the key projection matrix of B, and W_B^V is the value projection matrix of B.

Computing cross-attention for each pair of modalities incurs large computational costs and introduces redundancy in the fused features. As shown in Figure 6, we thus instead design the crossmodal Transformer to have the query statement as the main modality and to serve as the query in cross-attention, while other modalities serve as the keys and values in attention. The choice of the main modality is based on the insight that the SQL modality is directly related to a diagnosed slow query. Later, we report on experiments that validate this design.

Take the cross-modal attention between the query statement *S* and the execution plan *P* as an example and let E_S and E_P denote the embeddings of the two modalities. We capture cross-modal dependencies between these two and then extract the cross-modal feature h_P from the plan modality:

$$Q_S = E_S W_S^Q, K_P = E_P W_P^K, V_P = E_P W_P^V, h_P = \operatorname{softmax}(\frac{Q_S K_P^I}{\sqrt{d}}) V_P,$$
(7)

where W_S^Q , W_P^K , $W_P^V \in \mathbb{R}^{d \times d}$ are learnable feature projection and d is the dimension of the input embeddings. Applying similar calculations, we can obtain the cross-attention features from all the modalities: h_S , h_P , h_L , h_I . Having extracted the multimodal features that capture cross-modal relationships, the cross-modal Transformer (CMT) finally concatenates these along the feature dimension and obtains the fused representation through a feed-forward network (FFN):

$$CMT(E_S, E_P, E_L, E_I) = FFN([h_S; h_P; h_L; h_I])$$
(8)

Root-Cause-Adaptive Fusion. Features from multimodal fusion give a more comprehensive view of the slow queries, which helps the more accurate diagnosis of their root causes. In practice, as the query becomes more complex (e.g., with longer query statements, more nodes in the execution plan, larger numbers of returned rows, and longer execution time), only some parts of the information are highly related to the slow query and should be focused on in diagnosis. Furthermore, different root causes may have diverse relationships with different parts of multiple modalities. For example, when the root cause is a bad index, more attention should be given to the number of rows in the indexed field and the indexing method. This inspires us to fuse multimodal embeddings adaptively for diagnosing distinct root causes, as illustrated in Figure 4.

To distinguish the commonality and specificity of multimodal fusion among root causes, we decompose the features into two parts. One part involves features related to the commonality of the root causes, which are shared by all root causes. The other part involves features related to the specificity of each root cause, which are determined by their unique characteristics. We employ the Common Cross-Modal Transformer (denoted as C-CMT) to extract the fused representation E_F^C of the common features:

$$E_F^C = \text{C-CMT}(E_S, E_P, E_L, E_I) \tag{9}$$

To extract specific features for each root cause, we employ an Adaptive Cross-Modal Transformer (denoted as A-CMT) that adaptively fuses multimodal embeddings based on distinct root causes. For each root cause RC_j , we propose a gating mechanism to control the embedding of each modality during fusion. Take the query statement E_S as an example: the gating unit G_S^j outputs the gating value as the weight and then selects more relevant information for fusion:

$$G_{S}^{J}(E_{S}) = \text{sigmoid}\left(\text{FC}(E_{S})\right) \odot E_{S},$$
 (10)

where sigmoid is the activation function, \odot is the element-wise product, and FC is the submodule that output the gating value; we use a linear neural network layer for simplicity. Embeddings controlled by the gating unit then go through the cross-modal Transformer to extract specific multimodal features of the root cause RC_j :

$$E_A^j = \operatorname{A-CMT}\left(G_S^j(E_S), G_P^j(E_P), G_L^j(E_L)G_I^j(E_I)\right).$$
(11)

The final adaptive-fusion feature E_F^j for estimating the root cause RC_j is computed as the sum of the common feature E_C and the root-cause-adaptive feature E_A^j :

$$E_F^j = E_C + E_A^j \tag{12}$$

4.3.3 Unified Identification and Ranking.

Estimation of root cause impacts. In order to identify and rank valid root causes according to their influence on the slow query in a unified framework, we propose a training objective for estimating the impact of each potential root cause.

Given the multimodal input data X_i , the input embedding module and the multimodal fusion module extract the multimodal feature $E_F = \left\{ E_F^j \right\}_{j=1}^r$ as detailed in Sections 4.3.1 and 4.3.2, where E_F^j is the adaptively extracted feature for the root cause RC_j and r is the total number of all potential root causes. With E_F^j as input, the root cause estimation module estimates the impact of each root cause RC_j as \hat{y}_{ij} :

$$\hat{y}_{ij} = \mathrm{MLP}(E_F^j) \tag{13}$$

During training, we minimize the mean squared error (MSE) between the estimated and the annotated impacts of all the potential root causes:

$$\mathcal{L}_{\text{pred}} = \sum_{j=1}^{r} \|y_{ij} - \hat{y}_{ij}\|_2^2 \tag{14}$$

Recall that we regard root causes as valid if their impact exceeds a threshold ϵ . Here, we use the same threshold ϵ to apply the estimated impacts to obtain valid root causes and then rank them based on the estimated impacts.

Impact-aware Regularization. To further improve the accuracy of root cause identification and ranking, we propose two impact-aware regularizations to combine with the MSE training loss to obtain the overall training loss of the framework. The identification of root causes may be influenced by noise in the estimated impact, especially noise near the threshold ϵ . Motivated by this, we propose the valid regularization:

$$\mathcal{L}_{\text{valid}} = \sum_{j=1}^{r} \max(0, \mathbb{I}(y_{ij} < \epsilon) \cdot (\hat{y}_{ij} - \epsilon) + \eta), \quad (15)$$

where η is a margin, $\mathbb{I}(y_{ij} < \epsilon)$ is an indicator function that returns 1 if $y_{ij} < \epsilon$ and -1, otherwise. This encourages the estimated impact to be larger or smaller than threshold ϵ by at least a margin of η , depending on whether or not the root cause is valid. This regularization enhances the ability to distinguish between valid and invalid impacts in root cause identification.

Optimizing the model with only the MSE loss may not ensure an accurate ranking of the root causes due to the imperfect fit and estimations. We propose an order regularization to constrain the rankings of root causes. With $\{z_{ij}\}_{j=1}^r$ denoting sorted root causes according to ground truth impacts y_{ij} and $\{\hat{z}_{ij}\}_{j=1}^r$ denoting sorted root causes according to estimated impacts \hat{y}_{ij} , the regularization is:

$$\mathcal{L}_{\text{order}} = \sum_{j=1}^{r-1} \max(0, \left((z_{ij} - z_{i(j+1)}) - (\hat{z}_{ij} - \hat{z}_{i(j+1)}) \right))$$
(16)

The regularization constrains estimated orders to at least follow the margins between the ground truth orders, thereby preserving the impact ranking of the root causes. For example, given two root causes that have adjacent rankings, if the distance $\hat{z}_{ij} - \hat{z}_{i(j+1)}$ between the estimated results is 6.5% and the ground truth distance $z_{ij} - z_{i(j+1)}$ is 9.6%, we must update the estimated impact to be closer to the ground truth distance. This is so because the estimated distance is less than the ground truth distance, which affects the estimation of important root causes and even changes the order of root causes. Conversely, if the distance $\hat{z}_{ij} - \hat{z}_{i(j+1)}$ between the estimated results is 10.1, which exceeds the ground truth distance $z_{ij} - z_{i(j+1)}$ of 9.6, we does not perform any adjustment as this does not lead to incorrect estimation of important root causes or their order. The overall training loss of the framework is formulated as follows.

$$\mathcal{L} = \mathcal{L}_{\text{pred}} + \lambda (\mathcal{L}_{\text{valid}} + \mathcal{L}_{\text{order}}), \tag{17}$$

where λ is a trade-off hyper-parameter between the estimation loss and the impact-aware regularization losses.

4.3.4 Inference. We design an estimation model to estimate root cause impacts and rank them to determine the key root causes. As shown in Algorithm 1, given a slow query X_i , they are processed individually through the pre-trained models ENC_S, ENC_P, ENC_L, and ENC_I to obtain E_S , E_P , E_L , and E_I . Using C-CMT, we extract the common features shared by each root cause. Then, we utilize a gating mechanism to capture the different parts of the query statement, execution plan, execution log, and KPIs that each root cause focuses on. For each root cause, we use A-CMT to extract its specific features. We combine the common and specific features of each root cause and employ a Linear layer to estimate the impact of each individual root cause. Finally, we sort all root causes by their impact y_{ij} and eliminate those that are below impact threshold ϵ to obtain a list of valid root causes and their impact.

5 EXPERIMENTS

5.1 Experimental Design

5.1.1 Datasets. We collect five datasets, Hologres1, Hologres2, TPC-DS, TPC-C, and TPC-H, with the slow query threshold δ set to 1 second and the valid root cause threshold ϵ set to 10%. We consider 5 types of root causes in Hologres1, Hologres2, and TPC-DS: 1. outdated statistical information; 2. under-optimized join order algorithm; 3. missing or redundant indexes; 4. inappropriate distribution keys; 5. poorly written queries. In TPC-C and TPC-H, we consider 10 types of root causes: 1. outdated statistical information; 2. under-optimized join order; 3. inappropriate distribution keys; 4. missing indexes; 5. redundant indexes; 6. repeatedly executing subqueries; 7. complex table joins; 8. updating an entire table; 9. inserting large data; 10. unknown root causes. We provide detailed statistics of the datasets in Table 2, including the total number of

Algorithm 1: Pseudo-code of RCRank

Pre-training:

Input: All queries used for pre-training $\{Q_i | i \in [N]\}$ Initialize the encoder \mathbb{E}_{ϕ} ;

 $\tilde{S}, \tilde{P}, \tilde{L} \leftarrow \text{Mask } S, P, L \text{ in } Q_i;$

for $i \leftarrow 1$ to N do

for $\tilde{T} \in {\{\tilde{S}, \tilde{P}, \tilde{L}\}}$ do

Encode \tilde{T} and other original inputs using \mathbb{E}_{ϕ} ;

Reconstruct *I* in Q_i using encoder \mathbb{E}_{ϕ} ;

Update ϕ based on Eqs. 4 and 5;

Training:

Input: A root cause impact dataset $D = \{(y_{ij}, X_i, RC_j) | i \in [M], j \in [r]\}$ Initialize the weights θ of the whole estimation model with its encoder's weights ϕ initialized as the pre-trained weights;

for $i \leftarrow 1$ to M do

 $\begin{array}{l} E_{S}, E_{P}, E_{L}, E_{I} \leftarrow \text{Encode } S, P, L, I \text{ in } X_{i}; \\ E_{F}^{C} \leftarrow \text{Extract common features based on Eq. 9}; \\ \textbf{for } j \leftarrow 1 \textbf{ to } r \textbf{ do} \\ & \\ & \\ & \\ FuseFeatures(E_{S}, E_{P}, E_{L}, E_{I}, E_{F}^{C}); \\ & \\ & \\ \hat{y}_{ij} \leftarrow \text{Estimate impact of root cause } RC_{j}; \end{array}$

Update $\theta \leftarrow \arg \min_{\theta} (\mathcal{L}(y_{ij}, \hat{y}_{ij}))$ using Eq. 17;

Inference:

Input: A new slow query X_i $E_S, E_P, E_L, E_I \leftarrow$ Encode S, P, L, I in X_i ; $E_F^C \leftarrow$ Extract common features based on Eq. 9; **for** $j \leftarrow 1$ **to** r **do**

FuseFeatures($E_S, E_P, E_L, E_I, E_F^C$);

 $\hat{y}_{ij} \leftarrow \text{Estimate impact of root cause } RC_j;$

Return: Validate the root cause using \hat{y}_{ij} and sort \hat{y}_{ij} in descending order

Function FuseFeatures $(E_S, E_P, E_L, E_I, E_F^C)$:

Calculate the gate value using Eq. 10;

Extract specific features based on Eq. 11;

Return: Fuse the common and specific features using

Eq. 12;

queries (#Queries) that we monitored, the number of slow queries (#Slow Queries), and the number of tables (#Tables).

(1) **Hologres 1 & 2**: We collect two datasets of slow queries from Alibaba's production cloud database system Hologres, which contains a large number of real-world queries. The root cause distribution in dataset **Hologres1** tends to lean towards lacking indexes, while the root cause distribution in **Hologres2** is uniform. The slow queries come from routine executions of Alibaba's internal activities, involving databases with millions of records and complex conditions. We employ the rule-based and LLM-based methods outlined in Section 4.2 to collect the impacts of root causes.

Table 2: Details of datasets.

Dataset	#Queries	#Slow Queries	#Tables
Hologres1	12631	2209	136
Hologres2	13544	2701	84
TPC-DS	12196	7246	24
TPC-C	9980	5764	10
TPC-H	7638	4279	8

(2) TPC-DS, TPC-C, and TPC-H: These are standard benchmarks for evaluating the performance of database management systems. TPC-DS simulates a retail enterprise's decision support environment, encompassing complex queries, data mining, and business intelligence functionalities. TPC-C simulates the activities of a wholesale supplier. TPC-H is a decision support benchmark designed for the retail industry. It is possible to generate databases, tables, and specific mock data of chosen sizes. Although TPC-DS, TPC-C, and TPC-H offer query generation, only about a hundred generation templates is insufficient for our experimental needs. We thus use LLMs to generate more complex queries based on the relevant table structures in Postgres. Specifically, we use table structures such as table names and column names as prompts and require the LLM to generate complex query statements. After a selection process, we use the revision methods from Section 4.2 to perform root cause annotation on the slow queries, thereby producing synthetic datasets.

5.1.2 Baseline Approaches. (1) OpenGauss [15]: The SQL-level Diagnosis component in OpenGauss is similar to our task, which uses execution plans, and key performance indicators as input to find a time-consuming SQL operator. However, OpenGauss cannot rank the root causes. To enable OpenGauss to estimate an impact of the root cause, we fuse its execution plan and KPI features and use a fully connected layer for impact estimation. (2) D-BOT [52]: D-BOT collects diagnostic files and utilizes an LLM to automatically extract knowledge from these to diagnose the root causes of slow queries. (3) Only-SQL, Only-Plan, Only-KPI, Only-Log: We use the embedding module of our proposal to learn representations of a single modality and then estimate the impact of the root cause through a fully connected layer. (4) Concat: We use the embedding modules of our proposal to learn the representation of each modality and then concatenate them and estimate the impact of the root cause through a fully connected layer.

5.1.3 Evaluation Metrics. We evaluate the effectiveness of the methods based according to four metrics: impact error, root cause accuracy, root cause order, and key root cause.

First, we evaluate whether the estimated root cause impact values are accurate compared to ground truth impact values. We compute MSE error for each estimated impact values and then report the mean and standard deviation of the estimated impact values in the whole test set.

Second, we evaluate the accuracy of the valid root causes (*V*-*ACC*) and the most important valid root causes (*Top1*-*ACC*). Recall that a root cause is valid if its impact value exceeds the root cause threshold ϵ . Specifically, *V*-*ACC* calculates the accuracy of valid root causes by computing the ratio of correctly estimated root

causes to the total number of root causes. *Top1-ACC* calculates the accuracy of the root cause with the highest impact value.

Third, we evaluate the root cause orders obtained by the estimated vs. the ground truth impacts. We employ the multi-cause accuracy (MC-ACC) and Kendall's Tau (Tau) metrics. MC-ACC measures the accuracy of estimating the list of valid root causes sorted by the impact of each root cause for individual queries. Specifically, $MC-ACC = \frac{Order_{RC}}{m}$, where $Order_{RC}$ is the number of lists of estimated valid root causes that are ordered consistently with the ground truth. Tau evaluates the performance in terms of order accuracy: $Tau(RC) = \frac{n_c - n_d}{\sqrt{(n_0 - n_a)(n_0 - n_b)}}$, where n_c is the number of concordant pairs where the estimated impact values are consistent with the actual values in their order; n_d is the number of discordant pairs. n_a and n_b represent the number of tied rankings in the ordered lists of estimated root causes and ground truth root causes, respectively; and $n_0 = \frac{n(n-1)}{2}$, where n is the length of the list. In practical applications, the negative impact of providing incorrect root causes exceeds the negative impact of under-estimated root causes

Finally, to measure the improvement the estimated top-1 root cause can bring, we use *Top1-IR* to evaluate the performance gain after revising queries according to the top-1 root cause. Specifically, *Top1-IR* calculates the impact value of the estimated top-1 root cause for each slow query on average. In the process of revising slow queries, users prioritize revising slow queries according to the root cause with the greatest impact, i.e., the top-1 root cause. Higher *Top1-IR* values indicate that the top-1 root cause yields a larger performance improvement. Thus, the higher the *Top1-IR*, the better.

In Tables 3, 4, 5, 6, and 9, \uparrow indicates that higher values represent better performance, while \downarrow indicates that lower values represent better performance.

5.1.4 Implementation. We split the root cause impact dataset D into 8:1:1 for training, validation, and testing. We use all queries, except the validation and testing part of the slow queries, for pretraining encoders. We conduct training on an NVIDIA GeForce RTX 3090 GPU with a batch size of 64 and train for 50 epochs. We employ bert-base-uncased [7] for encoding queries, Queryformer [49] for encoding plans, a multi-layer perceptron for encoding logs, and a 2DCNN for encoding KPIs. The bert-base-uncased uses pre-trained parameters and the multi-layer perceptron consists of three fully connected layers with 13, 64, and 32 dimensions for the three layers. In the Cross Transformer, we set the number of blocks to 3 and the dropout rate to 0.1. We utilize the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a learning rate of lr = 0.0003. The trade-off for the losses is set to $\lambda = 7$, with further considerations of the value of λ in Section 5.2.5.

5.2 Experimental Results

5.2.1 Overall Performance. Tables 3, 4, and 5 show a comparative analysis of the proposed method and the baseline methods across diverse evaluation metrics on two real datasets and a synthetic dataset. Overall, RCRank outperforms baseline methods at root cause identification and ranking, providing evidence that RCRank is effective. Specifically, we observe the following.

(1) The single-modal approach performs relatively well on Only-SQL and Only-Plan, indicating that the estimation of root cause impact is more dependent on SQL and Plan. However, the single-modal approach lacks other crucial information; thus, these methods do not have a comprehensive understanding of slow queries. The query statement contains table structure information, the execution plan includes details like execution order and operations, the execution log provides information on actual resource consumption during execution, and KPIs reflect the current state of the system.

(2) OpenGauss, based on the concatenation fusion of execution plan and KPIs, only integrates information from the execution plan and KPIs, lacking details from query statements and useful information from execution logs for identifying root causes. In contrast, our approach utilizes information from query statement, execution plan, execution logs, and KPIs, fusing features through a Cross-modal Transformer across these four modalities. This allows us to make use of additional information relevant to root causes. The table shows that paying attention to information such as the semantic content in query statements and actual execution information in execution logs can enhance the diagnosis of slow query root causes and the estimation of their impact.

(3) Due to the insensitivity of LLMs to numbers, D-Bot cannot estimate the impact of root causes. Therefore, in D-Bot, we focus on the accuracy of single root cause estimation. Due to privacy concerns, we could not use GPT as the underlying LLM on Hologres. Instead, we opt for the open-source vicuna-13B-1.5k model, which we deploy locally. In scenarios where the use of GPT is not possible, D-Bot exhibits lower accuracy at single root cause estimation. This indicates that for root cause determination, LLMs still require more knowledge and the support of powerful language models as a foundation. On the TPC-DS, we use GPT-3.5 for our large model. We find that the trained model performs better a identifying the root causes of slow queries.

(4) For the task of identifying and ranking root causes, we focus on the performance improvement of slow queries after applying the revision methods based on root causes. To reduce revision costs, we focus on whether addressing estimating the top-1 root cause can bring substantial improvement. We observe that RCRank outperforms the other methods on the *Top1-IR* metric, providing evidence that integrating query statements, execution plans, execution logs, and KPIs is effective at estimating root causes and impacts.

5.2.2 Ablation Studies. To study the effect of different components, we create four variants of RCRank and conduct several experiments on both the Hologres and Postgres to compare their performance:

(1) **w/o gate unit**: To assess the importance of the Root Cause Gating Unit a providing attention to different parts of the input for different root causes, we remove the gating unit component from the RCRank. As shown in Figure 6, we observe that when not using the gate unit, performance is the worst. This suggests that different root causes focus on different parts of the input, e.g., diagnosing different root causes focuses on partial tokens in the query statement, partial operations in the execution plan, partial values in the logs, and partial KPIs in the KPIs.

(2) **use concatenation**: To assess the importance of the Crossmodal Transformer at fusing different features, we replace the Crossmodal Transformer with concatenation. This causes a decrease in

Table 3: Overall root cause diagnosis results in Hologres.

-																
		Mod	ality			Hologres1				Hologres2						
Method	SQL	Plan	Log	KPIs	V-ACC ↑	Top1-ACC ↑	MSE±std↓	MC-ACC ↑	Tau ↑	Top1-IR ↑	V-ACC ↑	Top1-ACC ↑	MSE±std↓	MC-ACC ↑	Tau ↑	Top1-IR ↑
Concat	\checkmark	~	\checkmark	~	0.7482	0.5071	0.1847 ± 0.1813	0.2132	0.3074	0.1751	0.7254	0.4240	0.2274±0.1793	0.3253	0.3319	0.2364
Only-SQL	✓				0.6898	0.4067	0.3595 ± 0.1953	0.1955	0.2269	0.1632	0.6989	0.3696	0.1056 ± 0.1954	0.2659	0.1748	0.1508
Only_Plan		\checkmark			0.7340	0.4743	0.2236 ± 0.1899	0.2281	0.2429	0.1680	0.6973	0.3803	0.4166 ± 0.1987	0.2925	0.2059	0.1584
Only-Log			\checkmark		0.6666	0.2900	1.8796 ± 1.2394	0.1613	0.1636	0.1592	0.6883	0.3590	0.1747 ± 0.2045	0.2659	0.2165	0.1422
Only-KPI				\checkmark	0.7058	0.3921	0.3989 ± 0.1854	0.1773	0.2230	0.1557	0.6922	0.3723	0.1522 ± 0.1923	0.2686	0.2093	0.1410
D-BOT (Vicuna)	√	\checkmark	\checkmark	\checkmark	0.3769	-	-	-	-	-	0.3518	-	-	-	-	-
OpenGauss		\checkmark		\checkmark	0.7354	0.5032	0.4362 ± 0.1870	0.2019	0.2740	0.1789	0.7149	0.4069	0.2335 ± 0.1847	0.3058	0.2518	0.1919
RCRank	\checkmark	~	\checkmark	\checkmark	0.7628	0.5384	$0.1226 {\pm} 0.1752$	0.2371	0.3518	0.1988	0.7420	0.4335	$0.0896 {\pm} 0.1645$	0.3563	0.3611	0.2764
Improvement					1.95%	6.17%	33.62%	3.95%	14.44%	11.12%	2.29%	2.24%	15.15%	9.53%	8.80%	16.92%

Table 4: Overall root cause diagnosis results in TPC-C and TPC-H.

			TPC-C				TPC-H					
Method	V-ACC↑	Top1-ACC ↑	MSE±std ↓	MC-ACC ↑	Tau ↑	Top1-IR ↑	V-ACC↑	Top1-ACC ↑	MSE±std ↓	MC-ACC ↑	Tau ↑	Top1-IR↑
Concat	0.8340	0.5619	0.1814 ± 0.1438	0.4917	0.4722	0.1413	0.8277	0.5743	0.1943 ± 0.1855	0.4253	0.4944	0.1677
Only-SQL	0.7602	0.5281	0.2794 ± 0.1774	0.4464	0.3965	0.1203	0.7425	0.5399	0.2914 ± 0.2216	0.4587	0.3492	0.1375
Only_Plan	0.7808	0.5307	0.2681 ± 0.2349	0.4791	0.4177	0.1294	0.7714	0.5344	0.2731 ± 0.1841	0.4757	0.4430	0.1294
Only-Log	0.7422	0.4711	0.3105 ± 0.2847	0.4359	0.4207	0.1176	0.7235	0.4933	0.3142 ± 0.2763	0.4188	0.3241	0.1025
Only-KPI	0.8054	0.5402	0.2347 ± 0.1547	0.4224	0.4796	0.1227	0.7827	0.5277	0.2467 ± 0.2164	0.4579	0.4428	0.1103
D-BOT(GPT-3.5)	0.6503	-	-	-	-	-	0.6108	-	-	-	-	-
OpenGauss	0.8216	0.5416	0.2103 ± 0.1697	0.4715	0.5176	0.1374	0.8146	0.5633	0.2120 ± 0.1977	0.4187	0.4703	0.1522
RCRank	0.8611	0.6417	0.1221±0.1378	0.5247	0.5574	0.1564	0.8523	0.6512	0.1496±0.1672	0.5718	0.5125	0.1712
Improvement	3.25%	14.20%	32.69%	6.71%	7.69%	10.69%	2.97%	13.39%	23.01%	20.20%	3.66%	2.09%

Table 5: Overall root cause diagnosis results in TPC-DS.

			TPC-DS	3		
Method	V-ACC ↑	Top1-ACC ↑	MSE±std↓	MC-ACC ↑	Tau ↑	Top1-IR ↑
Concat	<u>0.7836</u>	0.5566	0.7133 ± 0.1732	0.4547	0.4873	0.1201
Only-SQL	0.7381	0.5241	0.6133 ± 0.2677	0.4216	0.3224	0.1005
Only_Plan	0.7644	0.5142	0.4162 ± 0.2115	0.4331	0.4574	0.1014
Only-Log	0.6984	0.4869	0.9899 ± 0.4122	0.3824	0.3163	0.0973
Only-KPI	0.7713	0.5074	0.5849 ± 0.2413	0.4032	0.4130	0.0989
D-BOT(GPT-3.5)	0.5813	-	-	-	-	-
OpenGauss	0.7683	0.5243	0.3147 ± 0.1472	0.3926	0.4241	0.1142
RCRank	0.8466	0.6321	0.1732 ± 0.1440	0.5431	0.5391	0.1477
Improvement	8.04%	13.56%	44.96%	19.44%	10.63%	22.98%

performance, showing that cross-modal Transformers are more effective at learning integrated features. Cross-fused features can better understand query information, aiding in determining the types and impacts of root causes.

(3) **using MSE loss**: To assess the effect of Impact-aware Regularization, we replace the training loss with the MSE loss. We find that \mathcal{L}_{order} and \mathcal{L}_{valid} play crucial roles in the ordering and accuracy of identifying and ranking the root causes of slow queries. Imposing constraints on the ordering and validity of root causes can enhance the identification of effective and critical root causes. The MSE loss fails to distinguish root causes near the validity threshold. Since the MSE loss only calculates the mean square error, it cannot ensure the correct order of root causes when estimation errors exist.

(4) **w/o pre-train**: To assess the effect of the alignment pretraining method, we compare the performance with and without pre-training. We observe that the pre-training yields an improvement at estimating root causes and their impacts, indicating that using additional queries and aligning query statements, execution plans, and execution logs using the masking method and capturing patterns of KPIs using reconstructing method effectively capture the features of queries.



Figure 7: Sensitivity to parameter λ on Hologres and TPC-DS.

5.2.3 *Efficiency.* Table 7 shows the training time (per epoch) and inference time (per slow query) of RCRank, OpenGauss, and D-Bot, where we report the average runtime across 5 runs. In the training phase, D-Bot does not require training, since it leverages external knowledge, tools, and the internal knowledge of pre-trained LLMs for diagnosis. Therefore, we only compare the training time of RCRank and OpenGauss. We find that RCRank and OpenGauss have similar training times. The additional multimodal inputs do not incur significant computational overhead but result in better estimating the impact of root causes. In the inference phase, the machine learning methods with RCRank and OpenGauss are significantly faster than D-Bot. This is because D-Bot uses its LLM to make multiple calls to its dialogue model for multi-agent joint diagnosis to generate root causes, while RCRank and OpenGauss only require one call to output root causes and their impacts.

5.2.4 Original vs. Improved Run Time of Slow Queries. Our goal is to improve query efficiency by reducing the runtimes of slow queries, rather than completely revising them into non-slow queries. As shown in Figure 8, in most cases, the execution times of the revised slow queries are shorter than those of the original slow

Table 6: Ablation studies on Hologres and TPC-DS.

Dataset	Hologres							TPC-DS				
Method	V-ACC ↑	Top1-ACC ↑	MSE ↓	MC-ACC ↑	Tau ↑	Top1-IR ↑	V-ACC ↑	Top1-ACC ↑	MSE ↓	MC-ACC ↑	Tau ↑	Top1-IR ↑
w/o gate unit	0.7328	0.4400	0.1876	0.2340	0.2688	0.1954	0.8141	0.5907	0.1864	0.4766	0.4998	0.1351
use concat	0.7355	0.4548	0.2342	0.2688	0.3025	0.1923	0.8154	0.5942	0.2126	0.5074	0.5043	0.1403
use MSE loss	0.7238	0.4533	0.2624	0.2441	0.2805	0.1712	0.8014	0.5981	0.2867	0.4871	0.5114	0.1389
w/o pre-train	0.7360	0.4746	0.1119	0.2718	0.3269	0.2264	0.8179	0.6014	0.1931	0.5266	0.5161	0.1412
pre-train	0.7465	0.4985	0.1164	0.2732	0.3303	0.2360	0.8466	0.6321	0.1732	0.5431	0.5391	0.1477

Table 7: Training and inference running times on Hologres.

	Training time (per epoch)	Inference time (per slow query)
RCRank	1.76 m	0.018 s
OpenGauss	1.44 m	0.015 s
D-Bot	-	5.17 m



Figure 8: Original vs. improved run time of slow queries in Hologres.

Database	Original run time (s)	Revised run time (s)	Improvement
Hologres1	1736.52	1226.68	29.36%
Hologres2	1297.86	901.49	30.54%
TPC-DS	1852.54	1363.66	26.39%
TPC-C	1501.74	1092.22	27.27%
TPC-H	1030.33	781.202	24.18%

Table 8: End-to-end run time improvement.

queries, to varying extents. As shown in Table 8, RCRank identifies root causes that lead to substantial speedups in both Hologres and TPC datasets.

5.2.5 Sensitivity Analysis. To study the impact of different λ on performance, we investigate the average performance of the Hologres datasets and the performance of the TPC-DS dataset with values of λ in [1, 3, 5, 7, 10]. As shown in Figure 7, the performance changes only little across different λ settings. Because \mathcal{L}_{order} constrains both the order of the root causes and the gap between root cause impacts, we obtain more accurate estimates of root cause impacts.

5.2.6 Main Modality Study. To investigate which data modality is most important for determining root cause impacts, we use SQL, Plan, Log, and KPIs as the main modalities, each fused with the

Table 9: Main modality studies.

Modality	V-ACC ↑	Top1-ACC↑	MSE ↓	MC-ACC ↑	Tau ↑	Top1-IR↑
SQL	0.7360	0.4746	0.1119	0.2718	0.3269	0.2264
Plan	0.7263	0.4680	0.1077	0.2601	0.2490	0.1831
Log	0.7088	0.4241	0.1401	0.2165	0.2216	0.1663
KPI	0.7178	0.4552	0.1357	0.2296	0.2305	0.1721

other modalities. We evaluate the performance of using different modalities as the main modality on Hologres across all metrics.

Table 9 shows that using SQL as the main modality yields the best performance, not only in *Top1-IR* but also across other metrics. The database instance information contained in the KPI's modality plays an auxiliary role in determining root causes and their impact. Using KPIs that contain less information about the queries themselves as the main modality. The performance of using the Plan modality as the main modality is very similar to using SQL as the main modality because execution plans contain the columns involved in the queries and their execution orders. However, it is worth noting that execution plans in Plan do not reflect the actual execution plan during query execution. SQL benefits from using the pre-trained text encoder BERT, although the encoder is not specifically trained using SQL statements. However, table and column names in queries often contain semantic information, which can be captured by the text encoder, thereby reflecting the characteristics of the query.

6 CONCLUSION

This paper aims to identify root cause types and to rank them according to their impact on slow queries. We propose a multimodal Ranking framework for the Root Causes of slow queries (RCRank), which formulates diagnosis as a multimodal machine learning problem, leveraging multimodal information from query statements, execution plans, execution logs, and key performance indicators. We propose a multimodal diagnosis model, which enables expressive embeddings of heterogeneous modalities with modal alignment and task relevance, effective and adaptive fusion of multimodal features, and unified identification and ranking of root causes. Experimental results show that our method is capable of outperforming existing methods regarding root cause identification and ranking. In future research, it is of interest to study how to best make a revision plan w.r.t. a set of slow queries, and study how to efficiently incrementally [27] train RCRank to include new root causes.

ACKNOWLEDGMENTS

This work was partially supported by National Natural Science Foundation of China (62406112, 62372179) and Alibaba Innovative Research Program.

REFERENCES

- Mert Akdere, Ugur Çetintemel, Matteo Riondato, Eli Upfal, and Stanley B. Zdonik. 2012. Learning-based Query Performance Modeling and Prediction. In *ICDE*. 390–401.
- [2] Dana Van Aken, Andrew Pavlo, Geoffrey J. Gordon, and Bohan Zhang. 2017. Automatic Database Management System Tuning Through Large-scale Machine Learning. In SIGMOD. 1009–1024.
- [3] David Campos, Miao Zhang, Bin Yang, Tung Kieu, Chenjuan Guo, and Christian S. Jensen. 2023. LightTS: Lightweight Time Series Classification with Adaptive Ensemble Distillation. Proc. ACM Manag. Data 1, 2 (2023), 171:1–171:27.
- [4] Peng Chen, Yingying Zhang, Yunyao Cheng, Yang Shu, Yihang Wang, Qingsong Wen, Bin Yang, and Chenjuan Guo. 2024. Pathformer: Multi-scale Transformers with Adaptive Pathways for Time Series Forecasting. In *ICLR*.
- [5] Yunyao Cheng, Peng Chen, Chenjuan Guo, Kai Zhao, Qingsong Wen, Bin Yang, and Christian S. Jensen. 2023. Weakly Guided Adaptation for Robust Time Series Forecasting. Proc. VLDB Endow. 17, 4 (2023), 766–779.
- [6] Sudipto Das, Miroslav Grbic, Igor Ilic, Isidora Jovandic, Andrija Jovanovic, Vivek R. Narasayya, Miodrag Radulovic, Maja Stikic, Gaoxiang Xu, and Surajit Chaudhuri. 2019. Automatically Indexing Millions of Databases in Microsoft Azure SQL Database. In SIGMOD. 666–679.
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In NAACL-HLT. 4171–4186.
- [8] Bailu Ding, Sudipto Das, Ryan Marcus, Wentao Wu, Surajit Chaudhuri, and Vivek R. Narasayya. 2019. AI Meets AI: Leveraging Query Executions to Improve Index Recommendations. In SIGMOD. 1241–1258.
- [9] Dejan Dundjerski and Milo Tomasevic. 2022. Automatic Database Troubleshooting of Azure SQL Databases. *IEEE Trans. Cloud Comput.* 10, 3 (2022), 1604–1619.
- [10] Devamanyu Hazarika, Roger Zimmermann, and Soujanya Poria. 2020. MISA: Modality-Invariant and -Specific Representations for Multimodal Sentiment Analysis. In MM. 1122–1131.
- [11] Vimalkumar Jeyakumar, Omid Madani, Ali Parandeh, Ashutosh Kulshreshtha, Weifei Zeng, and Navindra Yadav. 2019. ExplainIt! - A Declarative Root-cause Analysis Engine for Time Series Data. In SIGMOD. 333–348.
- [12] Andreas Kipf, Thomas Kipf, Bernhard Radke, Viktor Leis, Peter A. Boncz, and Alfons Kemper. 2019. Learned Cardinalities: Estimating Correlated Joins with Deep Learning. In CIDR.
- [13] Tim Kraska, Mohammad Alizadeh, Alex Beutel, Ed H. Chi, Ani Kristo, Guillaume Leclerc, Samuel Madden, Hongzi Mao, and Vikram Nathan. 2019. SageDB: A Learned Database System. In CIDR.
- [14] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter A. Boncz, Alfons Kemper, and Thomas Neumann. 2015. How Good Are Query Optimizers, Really? Proc. VLDB Endow. 9, 3 (2015), 204–215.
- [15] Guoliang Li, Xuanhe Zhou, Ji Sun, Xiang Yu, Yue Han, Lianyuan Jin, Wenbo Li, Tianqing Wang, and Shifu Li. 2021. openGauss: An Autonomous Database System. Proc. VLDB Endow. 14, 12 (2021), 3028–3041.
- [16] Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. 2022. BLIP: Bootstrapping Language-Image Pre-training for Unified Vision-Language Understanding and Generation. In *ICML*. 12888–12900.
- [17] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. 2021. Align before Fuse: Vision and Language Representation Learning with Momentum Distillation. In *NeurIPS*. 9694–9705.
- [18] Greg Linden. 2006. Akamai online retail performance report: Milliseconds are critical. http://glinden.blogspot.com/2006/11/marissa-mayer-at-web-20.html
- [19] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. ACM Comput. Surv. 55, 9 (2023), 195:1–195:35.
- [20] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. In *NeurIPS*. 13–23.
- [21] Pengfei Luo, Tong Xu, Shiwei Wu, Chen Zhu, Linli Xu, and Enhong Chen. 2023. Multi-Grained Multimodal Interaction Network for Entity Linking. In SIGKDD. 1583–1594.
- [22] Lin Ma, Dana Van Aken, Ahmed Hefny, Gustavo Mezerhane, Andrew Pavlo, and Geoffrey J. Gordon. 2018. Query-based Workload Forecasting for Self-Driving Database Management Systems. In SIGMOD. 631–645.
- [23] Minghua Ma, Zheng Yin, Shenglin Zhang, Sheng Wang, Christopher Zheng, Xinhao Jiang, Hanwen Hu, Cheng Luo, Yilin Li, Nengjun Qiu, Feifei Li, Changcheng Chen, and Dan Pei. 2020. Diagnosing Root Causes of Intermittent Slow Queries in Large-Scale Cloud Databases. Proc. VLDB Endow. 13, 8 (2020), 1176–1189.
- [24] Ryan Marcus, Parimarjan Negi, Hongzi Mao, Nesime Tatbul, Mohammad Alizadeh, and Tim Kraska. 2021. Bao: Making Learned Query Optimization Practical. In SIGMOD. 1275–1288.
- [25] 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 (2019), 1705–1718.

- [26] Ryan Marcus and Olga Papaemmanouil. 2018. Deep Reinforcement Learning for Join Order Enumeration. In SIGMOD. 3:1–3:4.
- [27] Hao Miao, Yan Zhao, Chenjuan Guo, Bin Yang, Zheng Kai, Feiteng Huang, Jiandong Xie, and Christian S. Jensen. 2024. A Unified Replay-based Continuous Learning Framework for Spatio-Temporal Prediction on Streaming Data. *ICDE* (2024).
- [28] Vishvak Murahari, Dhruv Batra, Devi Parikh, and Abhishek Das. 2020. Large-Scale Pretraining for Visual Dialog: A Simple State-of-the-Art Baseline. In ECCV, Vol. 12363. 336–352.
- [29] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.
- [30] Zhicheng Pan, Yihang Wang, Yingying Zhang, Sean Bin Yang, Yunyao Cheng, Peng Chen, Chenjuan Guo, Qingsong Wen, Xiduo Tian, Yunliang Dou, Zhiqiang Zhou, Chengcheng Yang, Aoying Zhou, and Bin Yang. 2023. MagicScaler: Uncertainty-aware, Predictive Autoscaling. *Proc. VLDB Endow.* 16, 12 (2023), 3808–3821.
- [31] Xiangfei Qiu, Jilin Hu, Lekui Zhou, Xingjian Wu, Junyang Du, Buang Zhang, Chenjuan Guo, Aoying Zhou, Christian S. Jensen, Zhenli Sheng, and Bin Yang. 2024. TFB: Towards Comprehensive and Fair Benchmarking of Time Series Forecasting Methods. *Proc. VLDB Endow.* 17, 9 (2024), 2363–2377.
- [32] Rainer Schlosser, Jan Kossmann, and Martin Boissier. 2019. Efficient Scalable Multi-attribute Index Selection Using Recursive Strategies. In ICDE. 1238–1249.
- [33] Karl Schnaitter and Neoklis Polyzotis. 2012. Semi-Automatic Index Tuning: Keeping DBAs in the Loop. Proc. VLDB Endow. 5, 5 (2012), 478–489.
- [34] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. VideoBERT: A Joint Model for Video and Language Representation Learning. In *ICCV*. 7463–7472.
- [35] Ji Sun and Guoliang Li. 2019. An End-to-End Learning-based Cost Estimator. Proc. VLDB Endow. 13, 3 (2019), 307–319.
- [36] Xiu Tang, Sai Wu, Mingli Song, Shanshan Ying, Feifei Li, and Gang Chen. 2022. PreQR: Pre-training Representation for SQL Understanding. In SIGMOD. 204– 216.
- [37] Jindong Tian, Yuxuan Liang, Ronghui Xu, Peng Chen, Chenjuan Guo, Aoying Zhou, Lujia Pan, Zhongwen Rao, and Bin Yang. 2025. Air quality prediction with physics-guided dual neural odes in open systems. *ICLR* (2025).
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NeurIPS*. 5998–6008.
- [39] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *NeurIPS*.
- [40] Xinle Wu, Xingjian Wu, Bin Yang, Lekui Zhou, Chenjuan Guo, Xiangfei Qiu, Jilin Hu, Zhenli Sheng, and Christian S. Jensen. 2024. AutoCTS++: zero-shot joint neural architecture and hyperparameter search for correlated time series forecasting. VLDB J. 33, 5 (2024), 1743–1770.
- [41] Xinle Wu, Dalin Zhang, Miao Zhang, Chenjuan Guo, Bin Yang, and Christian S. Jensen. 2023. AutoCTS+: Joint Neural Architecture and Hyperparameter Search for Correlated Time Series Forecasting. Proc. ACM Manag. Data 1, 1 (2023), 97:1–97:26.
- [42] Sean Bin Yang, Chenjuan Guo, Jilin Hu, Jian Tang, and Bin Yang. 2021. Unsupervised Path Representation Learning with Curriculum Negative Sampling. In IJCAI. 3286–3292.
- [43] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. In *NeurIPS*.
- [44] Dong Young Yoon, Ning Niu, and Barzan Mozafari. 2016. DBSherlock: A Performance Diagnostic Tool for Transactional Databases. In SIGMOD. 1599–1614.
- [45] Wenmeng Yu, Hua Xu, Ziqi Yuan, and Jiele Wu. 2021. Learning Modality-Specific Representations with Self-Supervised Multi-Task Learning for Multimodal Sentiment Analysis. In AAAI. 10790–10797.
- [46] Yi Zhang, Mingyuan Chen, Jundong Shen, and Chongjun Wang. 2022. Tailor Versatile Multi-Modal Learning for Multi-Label Emotion Recognition. In AAAI. 9100–9108.
- [47] Chenyu Zhao, Minghua Ma, Zhenyu Zhong, Shenglin Zhang, Zhiyuan Tan, Xiao Xiong, LuLu Yu, Jiayi Feng, Yongqian Sun, Yuzhi Zhang, Dan Pei, Qingwei Lin, and Dongmei Zhang. 2023. Robust Multimodal Failure Detection for Microservice Systems. In SIGKDD. 5639–5649.
- [48] Kai Zhao, Chenjuan Guo, Yunyao Cheng, Peng Han, Miao Zhang, and Bin Yang. 2023. Multiple Time Series Forecasting with Dynamic Graph Modeling. Proc. VLDB Endow. 17, 4 (2023), 753–765.
- [49] Yue Zhao, Gao Cong, Jiachen Shi, and Chunyan Miao. 2022. QueryFormer: A Tree Transformer Model for Query Plan Representation. Proc. VLDB Endow. 15, 8 (2022), 1658–1670.
- [50] Renjie Zheng, Junkun Chen, Mingbo Ma, and Liang Huang. 2021. Fused Acoustic and Text Encoding for Multimodal Bilingual Pretraining and Speech Translation.

In ICML, Vol. 139. 12736-12746.

- [51] Xuanhe Zhou, Lianyuan Jin, Ji Sun, Xinyang Zhao, Xiang Yu, Shifu Li, Tianqing Wang, Kun Li, and Luyang Liu. 2021. DBMind: A Self-Driving Platform in openGauss. *Proc. VLDB Endow.* 14, 12 (2021), 2743–2746.
- [52] Xuanhe Zhou, Guoliang Li, Zhaoyan Sun, Zhiyuan Liu, Weize Chen, Jianming Wu, Jiesi Liu, Ruohang Feng, and Guoyang Zeng. 2024. D-Bot: Database Diagnosis System using Large Language Models. Proc. VLDB Endow. 17, 10 (2024), 2514– 2527.
- [53] Jun-Peng Zhu, Peng Cai, Boyan Niu, Zheming Ni, Kai Xu, Jiajun Huang, Jianwei Wan, Shengbo Ma, Bing Wang, Donghui Zhang, et al. 2024. Chat2Query: A

Zero-Shot Automatic Exploratory Data Analysis System with Large Language Models. In *ICDE*. 5429–5432.

- [54] Jun-Peng Zhu, Peng Cai, Kai Xu, Li Li, Yishen Sun, Shuai Zhou, Haihuang Su, Liu Tang, and Qi Liu. 2024. AutoTQA: Towards Autonomous Tabular Question Answering through Multi-Agent Large Language Models. *PVLDB* 17(12) (2024), 3920 – 3933.
 [55] Rong Zhu, Wei Chen, Bolin Ding, Xingguang Chen, Andreas Pfadler, Ziniu Wu,
- [55] Rong Zhu, Wei Chen, Bolin Ding, Xingguang Chen, Andreas Pfadler, Ziniu Wu, and Jingren Zhou. 2023. Lero: A Learning-to-Rank Query Optimizer. *Proc. VLDB Endow.* 16, 6 (2023), 1466–1479.