



Powering In-Database Dynamic Model Slicing for Structured Data Analytics

Lingze Zeng
National University of
Singapore
lingze@comp.nus.edu.sg

Naili Xing
National University of
Singapore
xingnl@comp.nus.edu.sg

Shaofeng Cai
National University of
Singapore
shaofeng@comp.nus.edu.sg

Gang Chen
Zhejiang University
cg@zju.edu.cn

Beng Chin Ooi
National University of
Singapore
ooibc@comp.nus.edu.sg

Jian Pei
Duke University
j.pei@duke.edu

Yuncheng Wu
Renmin University of
China
wuyuncheng@ruc.edu.cn

ABSTRACT

Relational database management systems (RDBMS) are widely used for the storage of structured data. To derive insights beyond statistical aggregation, we typically have to extract specific subdatasets from the database using conventional database operations, and then apply deep neural networks (DNN) training and inference on these subdatasets in a separate analytics system. The process can be prohibitively expensive, especially when there are various subdatasets extracted for different analytical purposes. This calls for efficient in-database support of advanced analytical methods.

In this paper, we introduce LEADS, a novel SQL-aware dynamic model slicing technique to customize models for specified SQL queries. LEADS improves the predictive modeling of structured data via the mixture of experts (MoE) and maintains efficiency by a SQL-aware gating network. At the core of LEADS is the construction of a general model with multiple expert sub-models trained over the database. The MoE scales up the modeling capacity, enhances effectiveness, and preserves efficiency by activating necessary experts via the SQL-aware gating network during inference. To support in-database analytics, we build an inference extension that integrates LEADS onto PostgreSQL. Our extensive experiments on real-world datasets demonstrate that LEADS consistently outperforms the baseline models, and the in-database inference extension delivers a considerable reduction in inference latency compared to traditional solutions.

PVLDB Reference Format:

Lingze Zeng, Naili Xing, Shaofeng Cai, Gang Chen, Beng Chin Ooi, Jian Pei, Yuncheng Wu. Powering In-Database Dynamic Model Slicing for Structured Data Analytics. PVLDB, 17(13): 4813 - 4826, 2024.
doi:10.14778/3704965.3704985

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/NLGithubWP/indices>.

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.
Proceedings of the VLDB Endowment, Vol. 17, No. 13 ISSN 2150-8097.
doi:10.14778/3704965.3704985

1 INTRODUCTION

Relational Database Management Systems (RDBMS) are extensively employed as the primary storage solution for structured data across various applications [27, 31, 38, 45]. They serve as a fundamental infrastructure for various domains and are critical to the operation of numerous businesses [22, 25, 62]. In the contemporary business landscape, structured data analytics via databases has become an indispensable component for driving business growth and success [22, 25, 41, 62]. Traditional structured data analytics approaches rely on database-driven filtering or aggregation operations to derive insights. However, these insights only offer a limited statistical view, which often fails to capture the complexity and intricacies of the underlying patterns [20, 44]. Fortunately, recent advancements in Deep Neural Networks (DNNs) open up new horizons for advanced analytics beyond simple statistical aggregation [10, 11, 18, 32].

At its core, exploiting DNNs for advanced structured data analytics comprises two main phases: training and inference [18]. The former primarily involves the construction of a DNN model and the training of this model on targeted data, while the latter utilizes the trained model to make predictions on new data. Notably, to deliver advanced DNN-driven analytics for informed decision-making, effectiveness, and efficiency are the two most important metrics to optimize for [9, 15, 29]. Specifically, effectiveness focuses on the inference phase, measuring the extent to which the predictions delivered by the model are accurate. Meanwhile, efficiency evaluates the requirements of the model in terms of *response time* and *computational resources* in both phases [29].

In real-world scenarios, analysts are often more interested in performing analytics on specific subsets of data. For instance, they may assess trends among patients diagnosed with a particular disease, or study behaviors of consumers in a certain age group. Consider the scenario illustrated in Figure 1, where an analyst examines the income situation within a specific demographic. This analysis informs critical decision-making, such as setting personal loan limits and adjusting local commodity prices. Naturally, the analyst seeks to build an effective predictive model, delivering accurate predictions for these subsets of tuples, meanwhile executing predictions efficiently with minimal response time and computational resources. However, there are two main challenges in achieving this objective.

First, achieving efficient training for effective predictive modeling across analyst-specified subdatasets is challenging. Conventionally, a single general model is trained to support inference across all data tuples [18, 22, 27]. This approach is efficient and requires training only one model. However, such a model, optimized to capture the common patterns and general behaviors of the whole dataset, is likely not as effective in providing accurate predictions as a dedicated model trained on a specific subdataset of interest. Taking the scenario in Figure 1 as an example, a model trained explicitly for the group of men living in NYC would probably identify finer-grained patterns and behaviors pertinent to this subdataset. Given sufficient training instances, this dedicated model could outperform the general model significantly. Nonetheless, training a separate model for each subdataset is computationally prohibitive due to the combinatorial nature of potential subdatasets.

Secondly, enhancing the efficiency of inference execution for subdatasets queried from the RDBMS is also challenging from a system perspective. A major challenge is the coordination of the RDBMS data processing and the execution of inference tasks. Many existing solutions support the inference process using two separate systems [14, 53, 60], which requires transferring the inference data from an RDBMS to the inference system. However, this process is time-consuming, susceptible to errors, and may violate privacy and security requirements [56]. More critically, additional data transfer overhead is introduced to adversely affect inference efficiency.

To address the above challenges, we propose a novel SQL-aware dynamic model Slicing (LEADS) technique, which makes use of the meta-information in SQL queries to dynamically customize a predictive model, deriving meaningful insights in analysis tasks. Specifically, LEADS first constructs a high-capacity general model consisting of multiple replicas of the base model. These replicas, termed as *experts* in MoE, are trained to specialize in different problem subspaces for effective predictive modeling. To enhance effectiveness via MoE without incurring reduced inference efficiency, LEADS incorporates a SQL-aware gating network, which generates sparse *gating weights* based on the SQL query embedding to select necessary experts. Such a *sliced model* is trained to tailor for SQL queries and dedicated to the specified subdataset for inference effectiveness while maintaining efficiency. To further improve efficiency, we build an extension to integrate the inference process within the database via User-Defined-Function (UDF) with three optimizations: efficient execution allocation, memory sharing, and state caching. These approaches obviate the data transfer in separate systems, reduce data copying overhead, and reduce the cost associated with frequent model loading. We summarize our main contributions as follows.

- We formulate the SQL-aware structured data analytics problem, which requires efficient and effective predictive modeling on subdatasets specified by SQL queries. To the best of our knowledge, this is the first work for solving the problem.
- We propose a novel SQL-aware dynamic model slicing technique LEADS, which scales up the modeling capacity by replicating the base model as experts and devises a SQL-aware gating network to dynamically customize models for analytics tasks on SQL-specified subdatasets.

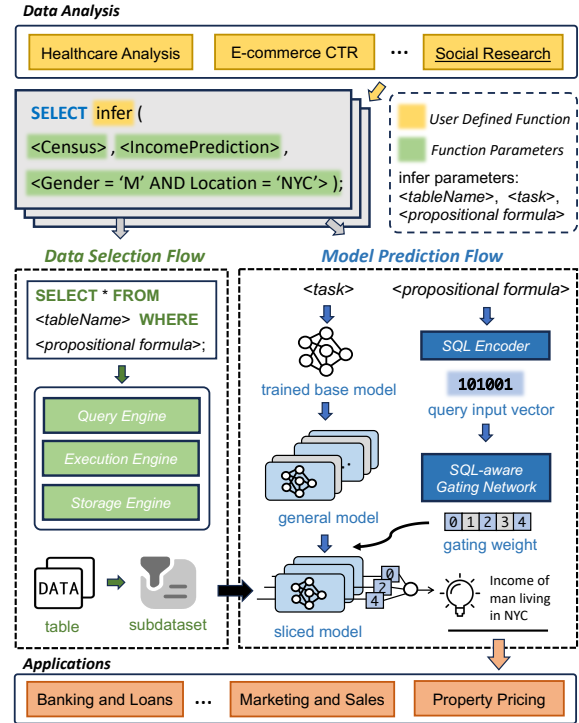


Figure 1: A workflow of supporting in-database analytics on income using SQL-aware dynamic model slicing.

- We implement an in-database inference extension on PostgreSQL to support SQL-aware structured data analytics, which incorporates three optimizations to improve inference efficiency.
- We conduct extensive experiments on five real-world datasets, which confirms the effectiveness of LEADS, with up to 3.95% improvement in AUC for given data workloads, while the inference extension achieves up to 2.06x speedup in terms of efficiency compared with the baseline approaches.

We have integrated LEADS with the extension onto NeurDB [39, 61], our ongoing AI-powered autonomous data system implementation.

The paper is structured as follows. We introduce preliminaries in Section 2. We present details of LEADS in Section 3 and discuss the implementation in Section 4. Experimental results are presented in Section 5. We review related work in Section 6 and conclude in Section 7.

2 PRELIMINARIES

In this section, we introduce the preliminaries of predictive modeling on structured data and present two key techniques central to our system, namely Mixture of Experts (MoE) for scaling up the model capacity while maintaining its inference efficiency via conditional computation [47], and sparsemax [35] for the informed selection of active experts to enhance efficiency. We also formally define the research problem, referred to as SQL-aware predictive modeling. Scalars, vectors, and matrices in the following part are denoted by x , \mathbf{x} and \mathbf{X} respectively.

Structured Data (or relational data, tabular data) [7, 8] refers to the type of data that can be represented as tables. It is typically stored in a series of tables (relations) $\{T_1, T_2, \dots\}$ of columns and rows, which can be retrieved from a relational database with SQL operations, e.g., the projection, join, and aggregation. Tables in structured data are interconnected via *foreign key* attributes, where the value of a column in one table points to a unique row in another table. Technically, the structured data can be viewed as a logical table \mathbf{T} comprising N rows and M columns. Each row is a tuple $\mathbf{x} = (x_1, x_2, \dots, x_M)$, serving as a feature vector, where x_i is the value of the i -th attribute and can be either numerical or categorical.

Mixture of Experts (MoE) [26, 47, 54] is a general ensemble learning and conditional computation technique to scale up the modeling capacity without incurring much computational overhead. MoE is particularly effective when the data exhibits complex patterns or variations [16, 46, 63]. There are two main components in an MoE layer: expert models and a gating network. Expert models can be composed of homogeneous models that share the same architecture. They are adopted to divide problem space into different regions, where each expert specializes in handling a certain sub-region. The gating network outputs a set of gating weights for a given input, dynamically determining weights assigned to respective experts.

Denoting the gating weights and outputs of experts as $\mathbf{w} = [w_1, w_2, \dots, w_K]$ and $\mathbf{H} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_K]$ respectively, where K is the number of experts and \mathbf{h}_i is the output of the i -th expert. The output of the MoE is a weighted average of these experts: $\hat{\mathbf{y}} = \sum_{i=1}^K w_i \mathbf{h}_i$. During training, the MoE model optimizes the gating network and experts simultaneously. The gating network learns to assign appropriate weight to experts, while the experts learn to make accurate predictions within their respective regions of expertise.

MoE has found extensive application in various domains, notably in the large language model GPT-4 [40] for texts and the large vision-language model MoE-LLaVA [33] for images, which combines the benefits of large model capacity with efficient computation, by only engaging a fraction of the model parameters for each input. In LEADS, we apply the MoE technique to structured data, intending to harness its scalable modeling capacity for enhanced predictive effectiveness and efficiency.

Sparsemax [10, 35, 42]. Softmax transformation is a crucial function in the gating network, which maps an input vector \mathbf{z} into a probability distribution \mathbf{p} whose probabilities correspond proportionally to the exponential of its input values, i.e., $\text{softmax}(z_j) = \frac{\exp(z_j)}{\sum_i \exp(z_i)}$. The output can be denoted as the class probabilities or weights indicating the importance of corresponding inputs.

The softmax function is extensively used in DNNs due to its differentiable and convex properties. However, the output probabilities of the softmax function are dense, leading to less interpretability and effectiveness [10, 19]. To overcome this limitation, α -entmax is proposed to generalize the dense and sparse softmax, offering a parameter to adjust the sparsity level of probability distribution. Specifically, given a d -dimension probability as $\Delta^d := \{\mathbf{p} \in \mathbb{R}^d : \mathbf{p} \geq 0, \|\mathbf{p}\|_1 = 1\}$, α -entmax [42] can be interpreted in the variational form with Tsallis α -entropies $H_\alpha(\mathbf{p})$ [50]:

$$\alpha\text{-entmax}(\mathbf{z}) = \underset{\mathbf{p} \in \Delta^d}{\operatorname{argmax}} \mathbf{p}^T \mathbf{z} + H_\alpha(\mathbf{p}) \quad (1)$$

where $H_\alpha(\mathbf{p}) = -\frac{1}{\alpha(\alpha-1)} \sum_j (p_j - p_j^\alpha)$ if $\alpha \neq 1$, else $H_1(\mathbf{p}) = -\sum_j p_j \log p_j$, the Shannon entropy. A larger α causes α -entmax to assign more zero in the probability distribution. Sparsemax is adopted in LEADS to sparsify the activated experts in MoE to enhance predictive efficiency.

Problem Formulation. In structured data analytics, data analysts typically focus on specific subdataset characterized by shared attributes. For example, analysts may assess the readmission rates among patients diagnosed with a certain disease, or predict the e-commerce click-through rate (CTR) within a particular age group. Typically, for advanced analytical tasks that involve prediction, WHERE statement in a SQL query is executed first to select relevant tuples, to which DNNs are applied subsequently for inference. We refer to this process as SQL-aware predictive modeling.

Given structured data \mathbf{T} as illustrated in Section 2 and a SQL query denoted by q , there are two main steps in SQL-aware predictive modeling: data selection and model prediction. Corresponding to relational algebra, a generalized SQL query selection q is expressed as $\sigma_\varphi(\mathbf{T})$, where σ is the unary operator for selection and φ is the propositional formula in q . Typically, φ consists of multiple predicates connected by logical operators. The selection $\sigma_\varphi(\mathbf{T})$ retrieves all tuples in table \mathbf{T} that satisfies φ , formally defined as $\sigma_\varphi(\mathbf{T}) = \{\mathbf{x} : \mathbf{x} \in \mathbf{T}, \varphi(\mathbf{x})\}$. For simplicity, the subdataset retrieved by the SQL query q is denoted as $\mathbf{T}_\varphi = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, where n is the number of tuples. Each tuple $\mathbf{x}_i \in \mathbb{R}^M$ comprises M attributes, and \mathbf{x}_i can be represented as a vector of features, i.e., $\mathbf{x}_i = [x_{i,1}, x_{i,2}, \dots, x_{i,M}]$. DNNs are then applied to perform prediction on these selected tuples, e.g., to predict the labels $\mathbf{y} = \{y_1, y_2, \dots, y_n\}$, aiming to derive meaningful insights, such as patients readmission rates in healthcare analytics or CTR in e-commerce. Technically, SQL-aware predictive modeling refers to making predictions on a selected subset of tuples retrieved from a logical table \mathbf{T} based on a SQL query q with a propositional formula φ .

3 SQL-AWARE DYNAMIC MODEL SLICING

In this section, we introduce our SQL-aware dynamic model slicing technique, LEADS, for supporting in-database predictive analytics within existing RDBMSs such as PostgreSQL. Unlike conventional machine learning approaches that make predictions based solely on features of individual instances, LEADS identifies common attributes of subdataset as specified by SQL queries and exploits this meta-information for customized, efficient and effective predictive modeling. For example, as illustrated in Figure 2b, the SQL query specifies instances sharing attributes of gender (male), age (24), and location (NYC). LEADS leverages these propositional formulas, i.e., φ in SQL queries, to customize the predictive model, thereby enhancing prediction accuracy for the specified subdataset.

We first propose a SQL query encoder to extract the information encapsulated in φ into a representation vector for subsequent predictive modeling. Next, we elaborate on the model architecture and data processing flow of LEADS for SQL-aware structured data modeling. Finally, we introduce the optimization schemes, with two novel regularization terms proposed to balance the effectiveness and efficiency of the modeling process.

3.1 SQL Query Encoder

In SQL-aware predictive modeling, the WHERE clause filters tuples based on a propositional formula φ . This formula consists of predicates, each imposing a condition on a specific attribute, like "gender = 'M'". These predicates are combined using logical operators such as "AND" or "OR" to form a propositional formula. Given the exponential number of possible predicate combinations, devising a general approach to effectively transform SQL queries into feature vectors for subsequent predictive modeling is challenging. To achieve this, we focus on individual queries, referred to as primitive SQL query. Considering a table T with M attributes, the j -th attribute denoted as A_j , each attribute is linked to either a numerical or categorical feature. Particularly, each numerical feature needs to be converted into a corresponding categorical feature through discretization, which will be detailed subsequently. In a primitive SQL query, each attribute A_j may be associated with zero or one predicate, with predicates across attributes conjoined using the logical operator \wedge (AND). Technically, a predicate for the attribute A_j in a primitive SQL query can be expressed as $P_j : A_j = a_j$, where $a_j \in \mathcal{D}_j \cup \{\Delta_j\}$, \mathcal{D}_j represents the domain of possible values for attribute A_j , and Δ_j denotes a default value assigned to A_j when it is not specified in the query. Figure 2a illustrates a valid primitive SQL query example, contrasting with two non-examples. Thus, the propositional formula φ can be represented as:

$$\varphi = P_1 \wedge P_2 \wedge \dots \wedge P_M.$$

The objective of the SQL query encoder is to generate a categorical feature vector \mathbf{q} for each primitive SQL query to represent the meta-information from φ , which is achieved by concatenating the attribute values of the predicates. Formally, the feature vector of the SQL query encoding can be obtained by:

$$\mathbf{q} = [q_1, q_2, \dots, q_M]$$


where q_j is the categorical attribute value for predicate P_j . Figure 2b demonstrates the transformation of a primitive SQL query into a feature vector. Notably, the numerical attribute "age" here is discretized before encoding, and any attribute A_i without predicates is assigned the default value Δ_i .

Discretization. Discretization is essential for encoding numerical attributes like weight or salary, as their infinite possible values make direct encoding infeasible. The discretization algorithm partitions the domain \mathcal{D} of each numerical attribute into a predetermined number of bins, and then, assigns values to their respective nearest bins. This process aims to retain the key information of numerical data in the embedding space for preserving predictive capacity.

To this end, we employ a supervised discretization approach that takes into account the correlation between numerical attributes and the predictive target. This method aims to maximize information value (IV), which measures the uncertainty reduction within each bin relative to the prediction target. Higher IV values indicate a significant decrease in uncertainty, thereby preserving the predictive capacity. Specifically, we introduce *OptBinning* [37] implementation for discretization to optimize IV while supporting constraints like the maximum bin count per attribute. The binning rule is learned on the training set and applied to the test set.

SELECT * FROM Census WHERE city = "NYC" OR city = "BOS" ⊗
 SELECT * FROM Census WHERE age > 25 AND gender = "Male" ⊗
 SELECT * FROM Census WHERE edu = "MSc." AND gender = "Male" ⊙

(a) Examples of a primitive SQL query.



(b) Process of encoding a SQL query.

Figure 2: SQL query encoder.

3.2 SQL-Aware Dynamic Model Slicing

The SQL feature vector \mathbf{q} , derived from the SQL query encoder, captures key information that can be exploited to tailor the predictive model for target subdatasets. Our objective is to construct a general model with sufficient modeling capacity and then customize this model based on \mathbf{q} to enhance efficiency and predictive performance in SQL-aware predictive modeling. While some approaches [34, 51, 52] modify the model's architecture or parameters to maintain its size, they can lead to unstable predictive performance. Conversely, ensemble approaches [17] improve performance by stacking models, but at the cost of additional computation. To balance effectiveness and efficiency, we introduce LEADS, which scales up the modeling capacity via MoE by replicating the base models as expert models, and integrates a SQL-aware gating network based on \mathbf{q} of SQL queries to selectively activate specific experts. The overview of LEADS is illustrated in Figure 3, and the key components following the data flow are introduced in this subsection.

Preprocessing Module There are two sets of input constructed for the SQL-aware prediction modeling given an input tuple. The first set of input is constructed for the gating network and can be uniformly represented as a categorical feature vector \mathbf{q} . $\mathbf{q} = [q_1, q_2, \dots, q_M]$ comprises M feature values from respective attribute fields, where numerical attributes need to be converted into categorical attributes via discretization, as described in the previous subsection. The second set is the attribute values of the input tuple $\mathbf{x} = [x_1, x_2, \dots, x_M]$, and each attribute value x_i can be either categorical or numerical.

For both \mathbf{q} and \mathbf{x} , each field of attribute value v_i (q_i/x_i) needs to be transformed into a corresponding embedding vector \mathbf{e}_i to participate in the subsequent predictive modeling. Specifically, each categorical attribute is transformed via *embedding lookup*, i.e., $\mathbf{e}_i = \mathbf{E}_i[q_i]$, $\mathbf{e}_i \in \mathbb{R}^{n_e}$, where n_e is the dimension of embedding, and \mathbf{E}_i is the embedding matrix of this categorical attribute. Note that different embedding vectors in \mathbf{E}_i correspond to their respective attribute values. As for each numerical attribute x_j in \mathbf{x} , the corresponding embedding vector is obtained by linearly scaling up a learnable embedding vector $\hat{\mathbf{e}}_j$, i.e., $\mathbf{e}_j = x_j \cdot \hat{\mathbf{e}}_j$. In this way, we can obtain fixed-size inputs, namely embedding vectors $\hat{\mathbf{q}} = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M]$ and $\hat{\mathbf{x}} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M]$.

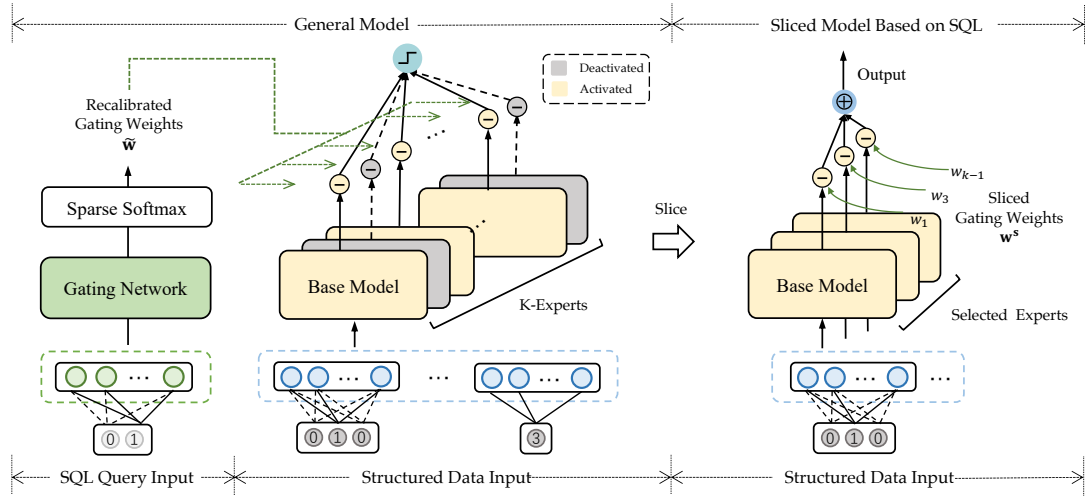


Figure 3: Overview of SQL-aware dynamic model slicing.

General Model and SQL-aware Gating Network. LEADS replicates the base model K times to construct the general model, denoted as $\mathcal{F} = [\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_K]$, where each base model is referred to as an “expert model” in MoE. These expert models share the same architecture but learn distinct model parameters during training, which take the same input $\hat{\mathbf{x}}$ and produce different outputs that are later aggregated to form the final prediction. The output of the i -th expert for a given input \mathbf{x} is denoted as $\mathcal{F}_i(\hat{\mathbf{x}})$.

The SQL-aware gating network \mathcal{G} takes the SQL query embedding vectors $\hat{\mathbf{q}}$ as input to produce a K -dimensional vector, termed the gating weight \mathbf{w} , with $\mathbf{w} \in \mathbb{R}^K$. Specifically, a two-layer multi-layer perceptron (MLP) is employed as the gating network following the practice [16, 40, 46]. We concatenate all embeddings in $\hat{\mathbf{q}}$ as the input of the gating network $\tilde{\mathbf{q}} = \mathbf{q}_1 \oplus \mathbf{q}_2 \dots \oplus \mathbf{q}_M$, where $\tilde{\mathbf{q}} \in \mathbb{R}^{M \cdot n_e}$, then feed $\tilde{\mathbf{q}}$ to \mathcal{G} , and obtain the gating weight \mathbf{w} by:

$$\begin{aligned} \mathbf{z} &= \phi(\mathbf{W}_1 \tilde{\mathbf{q}} + \mathbf{b}_1) \\ \mathbf{w} &= \mathcal{G}(\mathbf{z}) = \mathbf{W}_2 \mathbf{z} + \mathbf{b}_2 \end{aligned} \quad (2)$$

where $\mathbf{W}_1 \in \mathbb{R}^{n_z \times M n_e}$, $\mathbf{W}_2 \in \mathbb{R}^{K \times n_z}$ and $\mathbf{b}_1 \in \mathbb{R}^{n_z}$, $\mathbf{b}_2 \in \mathbb{R}^K$ are the weights and biases respectively, n_z is the hidden layer size, and ϕ represents the ReLU activation function.

Given the gating weight \mathbf{w} , the α -entmax function [13, 42] is further applied to recalibrate \mathbf{w} to a probability distribution. As introduced in Section 2, the hyper-parameter α in α -entmax controls the level of sparsity, and a larger value of α sets more gating weights to zero and thus deactivates more experts for higher efficiency. The output of α -entmax $\tilde{\mathbf{w}}$ is:

$$\tilde{\mathbf{w}} = \alpha\text{-entmax}(\mathbf{w}), \quad \tilde{\mathbf{w}} \in \mathbb{R}^K \quad (3)$$

which is utilized as weights to aggregate expert outputs. The final output of the general model is a weighted average of expert outputs:

$$\hat{\mathbf{y}} = \sum_{i=1}^K \tilde{w}_i \cdot \mathcal{F}_i(\mathbf{x}) \quad (4)$$

where $\hat{\mathbf{y}}$ is the prediction given the input \mathbf{x} and the corresponding query \mathbf{q} in the SQL-aware predictive modeling.

Dynamic Model Slicing via Gating Network. Given a SQL query, all retrieved data tuples share the same recalibrated gating weight $\tilde{\mathbf{w}}$. Notably, $\tilde{w}_i = 0$ in Equation 4 indicates that the i -th expert is not required for the current predictive task, and thus, only a selected subset of experts \mathcal{F}_i are activated, ensuring a more efficient inference process.

Denoting the set of indices of activated experts as $\{I_1, I_2, \dots, I_{n_o}\}$, where n_o is the current number of activated experts and $\tilde{w}_{I_j} \neq 0, \forall j \in \{1, 2, \dots, n_o\}$, and given the corresponding SQL feature vector \mathbf{q} , we index the activated experts to form a sliced model, i.e., $\mathcal{F}_q = [\mathcal{F}_{I_1}, \mathcal{F}_{I_2}, \dots, \mathcal{F}_{I_{n_o}}]$. Thus, the final output is as follows:

$$\hat{\mathbf{y}} = \sum_{j=1}^{n_o} \tilde{w}_{I_j} \cdot \mathcal{F}_{I_j}(\mathbf{x}) \quad (5)$$

where the number of activated experts n_o directly affects the effectiveness and efficiency of the sliced model. A large n_o indicates larger model capacity while incurring higher computational overhead, and vice versa. In LEADS, n_o is determined by the gating network based on the SQL feature vector \mathbf{q} and the hyper-parameter α of the sparsemax function. Notably, instead of predefining a fixed value, α in α -entmax is learnable and optimized based on the input tuples and corresponding queries during training. Subsequently, during inference, LEADS can dynamically adapt n_o based on the current SQL query, trading off between the effectiveness and efficiency of the predictive modeling.

3.3 Optimization

Our LEADS technique can be applied to different predictive tasks by configuring a proper objective function, such as mean squared error (MSE) for regression or cross-entropy for classification. For example, in binary classification, the objective function employed is binary cross-entropy:




Figure 4: Breakdown of inference response time for the Inference-Decouple Strategy.

$$\text{LogLoss}(\hat{y}, y) = -\frac{1}{N} \sum_i^N \{y_i \log \sigma(\hat{y}_i) + (1 - y_i) \log(1 - \sigma(\hat{y}_i))\} \quad (6)$$

where \hat{y} is the prediction label, y is the ground truth label, N is the number of tuples, and $\sigma(\cdot)$ is the sigmoid function.

To make the optimization more robust and effective, we introduce two regularization terms into the main loss function. The first term is the balance loss, \mathcal{L}_{baln} , which is to address *imbalanced expert utilization*, a common issue in training MoE-based models. This imbalance arises when the gating network \mathcal{G} disproportionately favor a few experts, which skews the training dynamics. As a consequence, preferred experts become overutilized while others remain underutilized, which compromises the MoE capacity and degrades overall model performance.

Let X denote a mini-batch of training instances with n_b tuples, and $\tilde{W} = [\tilde{w}_1, \tilde{w}_2, \dots, \tilde{w}_{n_b}]$ represent the recalibrated gating weights of X . Here, \tilde{w}_{ij} is the j -th weight of \tilde{w}_i . \mathcal{L}_{baln} is defined as follows:

$$\mathcal{L}_{baln} = cv(\Phi) = \sum_{j=1}^K \frac{\phi_j - \mathbb{E}(\Phi)}{\mathbb{E}(\Phi)^2} \quad (7)$$

$$\Phi = [\phi_1, \phi_2, \dots, \phi_K], \phi_j = \sum_{i=1}^{n_b} \tilde{w}_{ij}$$

where $\mathbb{E}(\Phi) = \frac{1}{K} \sum_{j=1}^K \phi_j$. This balance loss term promotes a uniform distribution of weights across experts in the mini-batch, encouraging equal contribution from each expert. However, this term empirically tends to activate a large number of experts due to its drive for balance, which is contrary to the intent of sparse softmax to maintain sparsity. To counterbalance this and enhance computational efficiency, we further introduce an additional term, the sparsity loss term, \mathcal{L}_{sprs} :

$$\mathcal{L}_{sprs} = -\frac{1}{n_b} \sum_{i=1}^{n_b} (\tilde{w}_i)^2 \quad (8)$$

which encourages the gating network to allocate higher weights on a select few experts while assigning minimal or zero weights to others. Both loss terms are weighted by their respective *regularization coefficient*, λ_1 and λ_2 , and incorporated into the overall objective function:

$$\text{Loss}(\hat{y}, y) = \text{LogLoss}(\hat{y}, y) + \lambda_1 \mathcal{L}_{baln} + \lambda_2 \mathcal{L}_{sprs}. \quad (9)$$

With this objective function, LEADS can then be optimized effectively using gradient-based optimizers, e.g., SGD [5] or Adam [28].

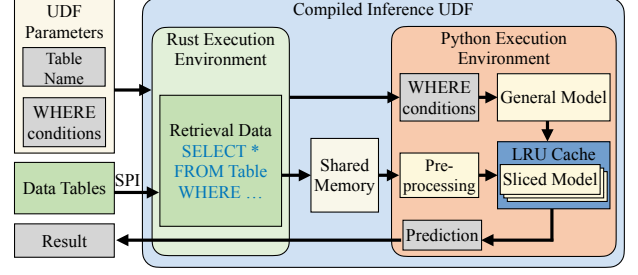


Figure 5: Design and execution of inference UDF.

4 IN-DATABASE MODEL INFERENCE

This section presents the implementation details of the in-database inference extension, which integrates the LEADS technique into NeurDB to deliver more accurate predictions. The typical inference pipeline consists of four stages: (1) model loading, where the model is loaded into memory; (2) data retrieval, where subdatasets specified by SQL queries are fetched; (3) data preprocessing, which transforms the raw data from the RDBMS into the tensor format for model inputs; and (4) inference execution, where the model performs forward computation to generate predictions. A straightforward approach to support LEADS is to decouple the database and inference systems, referred to as the **Inference-Decouple Strategy (IDS)**. In this approach, analysts retrieve subdatasets from the database, preprocess them, and run inference in a dedicated external inference system. However, it has several drawbacks. First, exporting data from the database introduces security risks and may violate compliance regulations. Second, managing two separate systems complicates the analytics workflow, increases operational complexity, and imposes additional learning requirements. Third, transferring large volumes of data from the database to an external system incurs significant overhead and latency. Figure 4 illustrates the time usage breakdown of IDS when running inference tasks. Notably, data retrieval accounts for around 40% of the total inference time primarily due to the overhead from database connections, data I/O, and network communication. To address these issues, we utilize the **In-database Inference Strategy (IIS)**, in which the database performs inference and avoids transferring data. This approach not only enhances security and simplifies workflow, but also accelerates the inference process.

4.1 Inference UDF Design

For the implementation, we design UDFs to encapsulate all stages of the inference process and integrate them into the database as an extension. As illustrated in Figure 5., users begin by initiating predictive queries, specifying parameters ‘TableName’ and ‘WHERE’ conditions. The UDF then retrieves relevant subdatasets by applying the ‘WHERE’ condition on specified table, dynamically loads the trained model, and performs the inference. Upon completion, the results are returned to the users. However, directly embedding the decoupled inference process into UDFs is suboptimal due to Python’s performance limitation. Therefore, we introduce three key optimizations to improve UDF execution efficiency.

Efficient Execution Allocation. To optimize UDF performance, we adopt a multi-language strategy, combining the strengths of both

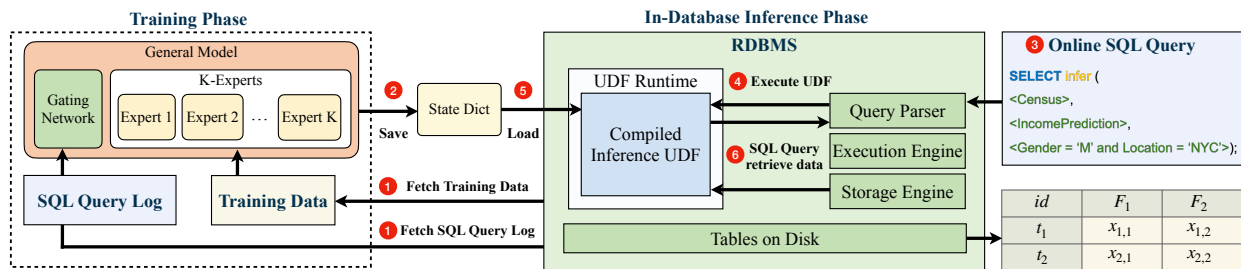


Figure 6: Workflow of in-database inference extension.

Rust and Python. Rust is used for efficient data retrieval, leveraging the PostgreSQL extension development library PGRX [2]. This allows us to utilize PostgreSQL’s built-in APIs, the Server Programming Interface (SPI) for high-performance data access. Meanwhile, Python is employed for data preprocessing and inference, benefiting from its compatibility with machine learning libraries like PyTorch and Scikit-learn. This combination ensures that we balance efficiency in data retrieval with the flexibility required for advanced data processing and machine learning tasks. However, two main challenges remain for efficient inference: (1) data copying overhead, caused by the need for extensive data transfers between different execution environments. (2) state initialization overhead, resulting from the frequent loading of predictive models for each inference task. To address these issues, we introduce memory sharing and state caching in our UDF designed to improve efficiency.

Memory Sharing. Data transfer between different execution environments requires two read-write operations: fetching data from RDBMS into Rust, then moving it to Python. To avoid this inefficiency, we introduce shared memory. Data is filtered in Rust using SPI and written directly to the shared memory, which is allocated before UDF invocation and accessible in both execution environment. The shared memory allows data to be used in Python without additional copying, eliminating unnecessary overhead.

State Caching. As shown in Figure 4, model loading takes up about 10% of the total inference time per request. With many requests, the overhead of repeatedly loading the model grows. To reduce it, we implement session-level caching for the general model and maintain a state cache for the frequently used sliced models. Specifically, when processing a SQL query with a new filter condition, the UDF first checks the state cache for an existing sliced model. If none is found, a new sliced model is derived from the cached general model. We employ a Least Recently Used (LRU) caching strategy to manage the model cache, ensuring efficient memory use.

4.2 System Workflow

In this section, we describe the workflow of the in-database inference extension, which comprises both training and in-database inference phases as shown in Figure 6.

Training Phase. In the training phase, a general model is constructed according to the table schemas and the predictive task SQL query logs from real-world scenarios are collected as the representative workload for training. These queries are utilized to retrieve the corresponding subdatasets, both of which are then preprocessed

and fed into the general model for iterative training (Step 1 in Figure 6). Upon completion of training, the general model is serialized and stored as a state dictionary (Step 2). In response to the incoming SQL query, an associated UDF is invoked, and then the model is loaded into the database memory and sliced to handle various online inference requests.

In-Database Inference Phase. The inference process is encapsulated into a UDF, which is integrated into database through extension installation. The UDF, named `infer`, provides a SQL interface for inference queries, as shown in the statement below:

```
SELECT infer(<tableName>, <task>, <propositional formula>);
```

Here, `<tableName>` indicates the source table for the subdataset. `<task>` specifies the prediction target (such as click-through rate or readmission rate) and directs to a particular deep learning model. `<propositional formula>` defines the filter conditions following the WHERE clause. Upon receiving the inference query (Step 3 in Figure 6), the UDF is executed as illustrated in Figure 5 (Step 4). The UDF performs four main tasks for an online inference request: (1) The prediction target and model are specified by `<tableName>` and `<task>`. If the model is not cached, the UDF will locate and load the general model (Step 5). (2) The UDF slices this model according to the `<propositional formula>` under the guidance of the LEADS technique. (3) Predictive data is retrieved via the SPI and written into shared memory (Step 6). (4) The sliced model is applied to the data for inference. Finally, the UDF returns a view containing the original data alongside a new predictive column.

5 EXPERIMENTS

In this section, we evaluate the effectiveness of LEADS and efficiency of our in-database inference extension, using five real-world datasets. We first introduce the datasets and experimental setup, and then design experiments and report findings to address four key research questions (RQs):

- RQ1: Does the LEADS technique improve the SQL-aware predictive modeling task compared to base models?
- RQ2: How effective is each component in LEADS technique?
- RQ3: Does our in-database inference extension improve the efficiency compared to traditional approaches?
- RQ4: How effective is our extension in complex scenarios?

Table 1: Dataset statistics.

Dataset	Tuples	Positive Ratio	Attributes	Features
Payment	30,000	21.4%	23	350
Credit	244,280	7.8%	69	550
Census	269,356	6.4%	41	540
Diabetes	101,766	46.8%	48	850
Avazu	40,428,967	17.2%	22	1,544,250

5.1 Experimental Setup

5.1.1 Datasets. We conduct experiments on five real-world datasets from the domains of finance, sociology, and healthcare. The statistics of the datasets are summarized in Table 1.

(1) **Payment** [58, 59] consists of credit card clients’ profiles and their past bill payments. The task is to predict if a credit card payment will default next month.

(2) **Credit** [6, 24] is gathered by Home Credit Group, focusing on the unbanked population. The task is to predict the repayment abilities of this population for better loan experience.

(3) **Census** [1, 57] contains data from the Current Population Survey conducted by the U.S. Census Bureau. The task is to determine whether a person’s annual income exceeds 50K based on their profile information, including age, class education, etc.

(4) **Diabetes** [12, 49] comprises ten years of clinical care at 130 US hospitals. Each tuple pertains to hospital records of patients diagnosed with diabetes, including details like medications and laboratory results. The task is to predict the patient’s readmission.

(5) **Avazu** [48] collects data from a mobile platform to predict click-through rates on ads. The dataset comprises millions of records and 22 attribute fields, covering mobile application and device information, with a total of 1,544,250 unique features.

5.1.2 Workloads. In SQL-aware predictive modeling, traditional OLAP benchmarks like TPC-DS [36] and TPC-H focus on evaluating query performance using operations such as JOIN and GROUPBY, which lack prediction tasks. In contrast, existing prediction datasets used for evaluating predictive modeling performance do not incorporate SQL queries. To address this deficiency, we develop a method to sample synthetic *inference queries* for the evaluation of SQL-aware predictive modeling.

Our method, detailed in Algorithm 1, employs a random strategy to create a set of synthetic SQL queries. The generation procedure is as follows: (1) randomly select a data tuple \mathbf{x} from the dataset D (Step 3); (2) sample a value m as the number of predicates in the current query from the range $[1, \min(\max_col, M)]$, where \max_col is the maximum number of predicates allowed in SQL queries, and M is the number of attributes in D (Step 4); (3) sample m attributes from \mathbf{x} and use their values to construct a propositional formula for the SQL query (Steps 5-6); (4) add the generated SQL query to the workload (Step 7). This procedure is repeated N to create a comprehensive workload. In experiments, we set N to 50 and \max_col to 3, and generate workloads on each dataset accordingly.

5.1.3 Baseline Methods. We select four kinds of base predictive models and enhance these models via the LEADS technique. We evaluate LEADS’s effectiveness by comparing the performance of these base models with and without the integration of LEADS. The base models are as follows

Algorithm 1 Synthetic Workload Generation

Require: dataset D , the number of SELECT queries N , the maximum filter condition size \max_col

Ensure: a synthetic workload W containing N SELECT queries

```

1:  $W = \emptyset$ 
2: for  $i \leftarrow 1$  to  $N$  do
3:   Randomly select a data tuple  $\mathbf{x} \in \mathbb{R}^M$  from  $D$ 
4:   Randomly sample the number of selected columns  $m \in [1, \min(\max\_col, M)]$ 
5:   Randomly sample  $m$  columns from data tuple  $\mathbf{x}$  along with their corresponding values
6:   Form a SELECT query with a filter condition of size  $m$  based on the selected columns and values
7:   Add the generated SELECT query to the workload  $W$ 
8: end for
9: return synthetic workload  $W$ 

```

(1) **DNN** [18]: contains fully-connected linear and activation layers, representing the most fundamental neural network.

(2) **CIN** [32]: it models higher-order feature combinations through compressed interaction with input embeddings.

(3) **AFN** [11]: it incorporates logarithm neurons in the network layer, aiding in capturing the feature interaction in arbitrary order.

(4) **ARMNet** [10]: it introduces multi-head attention to adaptively extract the combination of features, demonstrating state-of-the-art performance in structured data prediction tasks.

Additionally, to evaluate the efficiency of the in-database inference extension, we compare it with the traditional analytic approach IDS described in Section 4, in which data is retrieved out of the database through network communication based on *psycopg*, while no data is copied between execution environments.

5.1.4 Evaluation Metric. A workload consists of a set of inference queries, each representing a prediction request. We evaluate the effectiveness of LEADS on a single inference query, denoted by q , using the AUC (Area Under the ROC Curve) metric, where higher values indicates better performance. To assess LEADS’s overall performance across the entire workload, we use two metrics. The first is the average AUC for all queries, denoted as *Workload-AUC*:

$$\text{Workload-AUC}(W) = \frac{1}{N} \sum_{i=0}^N \text{AUC}(q_i) \quad (10)$$

where N is the number of SQL queries, q_i is the i -th query in the workload W . The second is the lowest AUC value among all queries, termed *Worst-AUC* and calculated as:

$$\text{Worst-AUC}(W) = \text{Min}\{\text{AUC}(q_1), \text{AUC}(q_2), \dots, \text{AUC}(q_N)\} \quad (11)$$

In fields like finance and healthcare where mistakes can lead to significant losses, it’s important to focus on the worst-case performance. The Worst-AUC reveals whether the technique performs reliably and helps avoid poor decisions. To assess model-level efficiency, we utilize the floating-point operations per second (*FLOPs*) to measure the computations during the inference phase. For system-level evaluation, we measure the performance using *response time*, which tracks the CPU time elapsed from when a user initiates a query to when the prediction results are received.

Table 2: Evaluation of performance improvements with LEADS.

Datasets	Metric	DNN			CIN			AFN			ARMNet		
		w/o	w/	Imprv.	w/o	w/	Imprv.	w/o	w/	Imprv.	w/o	w/	Imprv.
Payment	Workload-AUC	0.7003	0.7089	+1.23%	0.7164	0.7189	+0.35%	0.7067	0.7143	+1.08%	0.7141	0.7212	+0.99%
	Worst-AUC	0.4733	0.5467	+15.51%	0.3836	0.4463	+16.35%	0.4467	0.6333	+41.77%	0.5267	0.6067	+15.19%
Credit	Workload-AUC	0.7145	0.7427	+3.95%	0.7234	0.7408	+2.41%	0.7171	0.7218	+0.66%	0.7231	0.7347	+1.60%
	Worst-AUC	0.3852	0.6000	+55.76%	0.3333	0.4074	+22.23%	0.3852	0.4074	+5.76%	0.4444	0.6264	+40.95%
Census	Workload-AUC	0.9157	0.9200	+0.47%	0.9187	0.9224	+0.40%	0.9151	0.9216	+0.71%	0.9196	0.9237	+0.45%
	Worst-AUC	0.7692	0.8041	+4.54%	0.7692	0.7845	+1.99%	0.7577	0.7892	+4.16%	0.7692	0.7962	+3.51%
Diabetes	Workload-AUC	0.8308	0.8375	+0.81%	0.8322	0.8419	+1.17%	0.8329	0.8390	+0.73%	0.8342	0.8402	+0.72%
	Worst-AUC	0.5495	0.6374	+16.00%	0.6264	0.7033	+12.28%	0.6484	0.6813	+5.07%	0.6044	0.6593	+9.08%
Avazu	Workload-AUC	0.7355	0.7424	+0.94%	0.7324	0.7443	+1.62%	0.7364	0.7424	+0.81%	0.7387	0.7440	+0.71%
	Worst-AUC	0.4231	0.5562	+31.46%	0.4531	0.5625	+24.14%	0.5031	0.5594	+11.19%	0.4615	0.5625	+21.88%




Figure 7: Top-4 SQL queries in terms of AUC improvement.

5.1.5 *Hyper-parameter Settings.* For fair comparisons, we set the feature embedding size to 10 and the hidden layer size to 32 for all base models. Since LEADS allows multiple experts, we reduce the hidden layer size of each expert to 16 for efficiency. For ARMNet, we set the self-attention module to 4 heads and a hidden size of 16. The initial α in sparsemax is 2.5. The number of experts in LEADS is selected between 2 and 256 and fixed at 16. The balance regularization factor λ_1 and sparsity factor λ_2 are chosen between $1e-3$ and $5e-2$, both fixed at $1e-3$. We conduct a sensitivity analysis on these hyperparameters and report the best results.

5.1.6 *Training Details.* Since the training dataset lacks specific SQL queries to update the gating network parameters, we simulate a SQL query for each input following Steps 3-6 in Algorithm 1. We use the Adam optimizer [28] with a learning rate range of $1e-3$ to 0.1 and a batch size of 1024 for all base models and datasets. Experiments are conducted on a server with a Xeon Silver 4114 CPU @ 2.2GHz (10 cores), 256GB of memory, and a GeForce RTX 3090 Ti. All models are implemented using PyTorch 1.6.0 with CUDA 10.2.

5.2 SQL-aware Predictions

To answer the question RQ1, we investigate the performance of four base models with LEADS, and the results are summarized in Table 2. The main observation is that the prediction performance w.r.t. *Workload-AUC* and *Worst-AUC* consistently improve when utilizing LEADS on base models. Notably, the most significant improvement is observed in the *Worst-AUC* metric. For instance, when using DNN as the base model, LEADS achieves improvements of 55.76% and 16.00% on the Credit and Diabetes datasets, respectively. The




Figure 8: Effects of SQL-aware gating network on accuracy.

Table 3: Effects of the number of predicates on AUC.

Dataset	#predicates	1	2	3	4	5	6
Credit	#tuple	24375	9902	9900	9896	9856	5208
	w/o LEADS	0.7432	0.7495	0.7495	0.7494	0.7483	0.7494
	w/ LEADS	0.7522	0.7589	0.7591	0.7581	0.7577	0.7604

reason for the base model’s low performance could be significant variability or nuances in the instances of the retrieved subset that are not well-represented in the training data. Consequently, the trained base model fails to provide accurate predictions.

To further analyze the *Worst-AUC* improvement, we perform a breakdown analysis on the Diabetes dataset with the DNN base model. Figure 7 presents the top-4 SQL queries with the highest AUC improvements from LEADS, along with their query profiles. We notice that these queries have lower AUC values compared to the overall *Workload-AUC* (see Table 2) and involve small subsets of the training data. For instance, query#1 uses only 0.13% of the training data, about 130 out of 101,766 tuples. With such limited data, the base model struggles to generalize, resulting in poor predictions. LEADS addresses this by leveraging the propositional formulas in SQL queries to help the general model identify patterns in these small subdatasets, enhancing performance on queries with few training samples. In Diabetes dataset, many queries involve the drug status. For example, query#3 indicates a reduced dosage of *glipizide*, which significantly impacts readmission rates. This demonstrates LEADS’s value in improving healthcare predictions.

We also investigate the impact of predicate numbers on predictive results, as it represents query complexity. In our experiment,




Figure 9: Effects of α -entmax and number of experts.




Figure 10: Effects of the regularization terms on accuracy.

we incrementally added up to 6 predicates to evaluate the prediction. As shown in Table 3, LEADS outperforms baseline methods across different levels of query complexity. Additionally, there is an upward trend in predictive performance with the increase of the predicate number, indicating that complex primitive queries provide richer meta-information and enhance prediction accuracy.

5.3 Ablation Study

In this part, we conduct the ablation study to answer the question RQ2, evaluating the effectiveness of each component in LEADS.

SQL-aware gating network. In this evaluation, we remove the *SQL-aware gating network* to demonstrate the importance of dynamic expert selection. A default model is created using a special SQL query embedding, where a set of padding values for each attribute, denoted as $\mathbf{q}_d = [\Delta_1, \Delta_2, \dots, \Delta_M]$, indicating the absence of predicates in the SQL query. The comparison results are shown in Figure 8. There are two main observations. First, the *w/o LEADS* method achieves the lowest Workload-AUC because it simply uses the base model to handle all SQL queries. Second, the *LEADS w/o SQL-aware gating network* method results in a performance reduction compared to LEADS. For example, on the Credit dataset, this reduction can reach up to 0.02 in terms of Workload-AUC. It is because without the SQL-aware gating network, LEADS loses the ability for dynamic model customization based on SQL query vectors, leading to unsatisfactory results.

α -entmax. We evaluate the effect of the α -entmax function by comparing LEADS to *w/o LEADS* and *LEADS w/o α -entmax* (which uses the softmax function) Using DNN as the base model, we vary the number of experts from 2 to 256 and measure the performance w.r.t. Workload-AUC and FLOPs on Payment and Credit datasets. As shown in Figures 9, increasing the number of experts from 2 to 256 leads to notable AUC improvements, as more experts allow the model to make more accurate predictions. However, beyond 32




Figure 11: Analysis of the activated frequency for each expert. The value denotes the frequency and 1.0 means the expert is activated in every SQL query.

experts, LEADS w/o α -entmax sees a decline in AUC on the Credit dataset due to easily overfitting, as the model becomes too complex. Figure 9 also highlights the advantages of α -entmax in terms of FLOPs saving. Specifically, the FLOPs of LEADS w/o α -entmax grow linearly with more experts, while LEADS with α -entmax has a gentler increase, because α -entmax assigns small values to zero, reducing the number of active experts compared to softmax. We remove these unused experts to conserve computational resources.

Regularization terms. We compare the performance of LEADS with three variants: without the balance term (*LEADS w/o L_{baln}*), without the sparsity term (*LEADS w/o L_{sprs}*), and without both terms (*LEADS w/o both*). Using DNN as the base model, we conduct experiments on the Payment and Credit datasets. Figure 10 presents the comparison results in w.r.t. Workload-AUC, and Figure 11 analyzes the frequency of expert usage during query execution. Three main findings emerge: First, removing the balance term significantly reduces the Workload-AUC of LEADS, Without the balance term, fewer experts are used for each query, leading to lower prediction accuracy. This is evident in Figure 11, where only two experts are predominately selected. Second, adding solely the balance term results in lower performance than LEADS but utilizes almost all experts for every query. The balance term encourages even expert selection, leading to higher computational costs. Lastly, enabling both terms simultaneously in LEADS balances expert usage while achieving the best performance.

5.4 System Efficiency

To answer question RQ3, we assess the system-level efficiency in end-to-end response time by comparing the IIS that is used in our system with IDS. We implement IIS as an extension that can be installed in PostgreSQL 14.

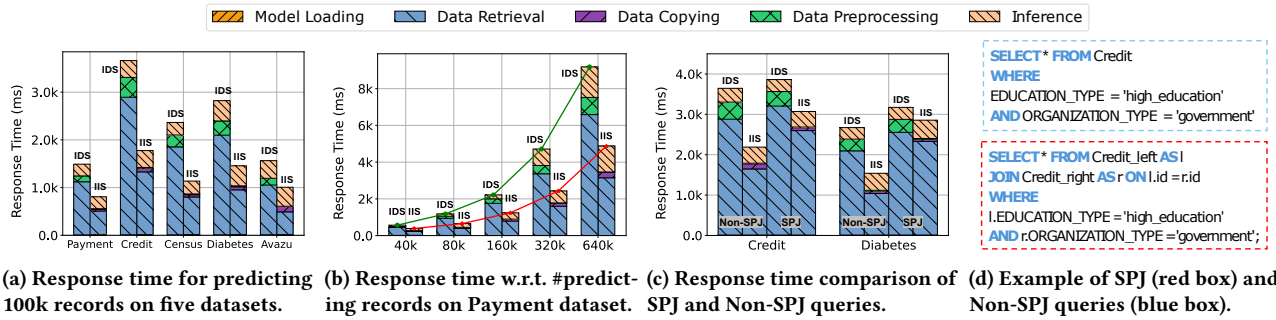


Figure 12: Efficiency evaluation of In-database Inference Strategy (IIS) and Inference-Decouple Strategy (IDS).




Figure 13: Effects of optimizations on response time.

Comparison with the baseline. Given a SQL query that selects 100k records for inference, We report the response time of IIS and IDS on the five datasets, as presented in Figure 12a. Compared to IDS, IIS achieves a speedup of 1.94x, 2.06x, 2.00x, 1.82x, and 1.53x on the Census, Credit, Diabetes, Payment, and Avazu. There are two reasons for such superior performance. One, IIS reduces the costly data movement overheads between PostgreSQL and the inference system, with lower data retrieval time usage. Secondly, IIS is further enhanced with the optimizations: shared memory to reduce data copying overhead, and state caching to eliminate the cost of model loading during the inference UDF execution process.

Effects of the number of predicting records. Next, we examine how the number of predictive records in the SQL query impacts response time. In the Payment dataset, this number ranges from 40k to 640k records. Figure 12b shows the response time for two strategies, IIS and IDS. We observe that IIS consistently surpasses IDS across various record numbers, with performance improvement ranging from 1.47x to 1.93x. Moreover, the response time of IIS increases more slowly than that of IDS, since the data movement overhead between the database and the inference system becomes more pronounced with the increase of records. Therefore, with more predictive records, IIS performs even more favorably.

Evaluation of optimization techniques Further, we evaluate the benefits of the optimizations mentioned in Section 4.1. Specifically, we compare the in-database inference process with (i) *w/o memory sharing*; (ii) *w/o SPI*; (iii) *w/o state caching*; and (iv) *w/o all optimizations*. Figure 13 presents the comparison results w.r.t. the response time in predicting 100k records on the Payment dataset. The absence of shared memory leads to significant overhead due to data copying between different execution environments. Likewise,




Figure 14: Optimizing UDFs: analyzing the trade-offs between response time and memory usage.

without SPI (PostgreSQL’s built-in data access API), data retrieval times are considerably longer. Moreover, there is considerable overhead from repeatedly loading the model without state caching. When all optimizations are enabled, the in-database inference extension achieves a 3x speed improvement compared to it without optimizations. Additionally, we investigate the trade-offs between response time and memory usage brought by memory sharing and state caching. In this analysis, we set a workload containing 300 inference queries and record memory usage during the process. The results presented in Figure 14 reveal that state caching and memory sharing enhance inference efficiency, with minimal impact on overall memory usage from these optimizations.

5.5 Complex Inference Scenarios

In this subsection, we address question RQ4 by examining LEADS’s effectiveness under data updates and schema changes, as well as the efficiency of our in-database inference extension when handling SPJ (Select-Project-Join) queries.

Effectiveness when data updates. We evaluate the effectiveness of LEADS for tuple insertion and deletion by executing the same SQL queries between data changes. Using the Credit and Diabetes datasets with a DNN as the base model, each dataset is equally divided into original and new tuples. The new tuples are inserted into the database in five stages, with the model evaluated after each insertion. As shown in Figure 15, the model enhanced by LEADS consistently outperforms the base model as the dynamic combination of multiple experts discussed in Section 5.2. In conclusion, when data updates follow the same distribution, the model enhanced by LEADS effectively generalizes to various subdatasets.




Figure 15: Effects of data insertion on Workload-AUC.

Effectiveness when the schema changes. As for schema changes, altering attributes changes the feature dimension of each tuple, leading to a dimension mismatch with the predictive model. To continue using the original model, we need to construct matched inputs. For insertion, required attributes can be fetched and fed to the model, while newly added attributes will not be considered in predictive modeling. For deletion, the values in the deleted attributes will be missing. We pad these missing values, but this will degrade performance. The impact of deleted attributes on model performance is shown in Figure 16. We notice that as more attributes are deleted, model performance drops rapidly. Therefore, retraining the model based on the new feature dimensions is advisable to ensure effective inference following schema changes.

Efficiency for SPJ queries. To evaluate in-database inference on SPJ queries, we vertically split the Diabetes and Credit datasets into two sub-tables (e.g., diabetes_left and diabetes_right). During inference, these sub-tables are joined to select tuples, as shown in the SQL example in Figure 12d, and the efficiency is compared to the baseline approach IDS. We also assess performance using the original table without JOIN to measure the impact of query complexity. As expected, Figure 12c shows that JOIN operations introduce additional overhead, leading to higher response times compared to queries without JOIN. However, IDS outperforms the baselines due to reduced data movement, as data remains within the database. For JOIN queries, the speed advantage of in-database inference is less pronounced since query planning and execution times dominate over the benefit of eliminating network latency.

6 RELATED WORK

Mixture-of-Experts (MoE) is initially proposed by [26] to handle different samples using independent expert modules. [47] introduces the Sparse Gated MoE into language model training, developing large-scale LSTM-based MoE models, in which only one expert is chosen for each input data. With the rise of the Transformer as the dominant NLP architecture, researchers began to incorporate MoE layers by extending the Feed-Forward Networks (FFNs) in Transformers to build MoE language models. Despite these innovations, the Sparse Gated MoE [47] struggled with stability due to the expert route strategy where it brute-force selects the top-1 expert. To mitigate this, various studies [16, 30, 43, 63] explored different learnable routing strategies for managing experts and input tokens. For example, Fast MoE [21] monitors the training status and dynamically adjusts the load for each expert. Expert Choice Routing [63] lets experts choose tokens rather than selecting experts to prevent under-training. Soft MoE [43] performs an implicit soft assignment

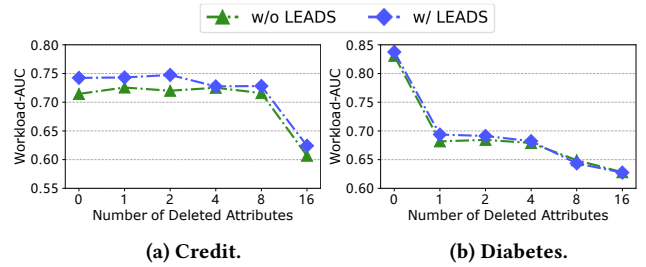


Figure 16: Effects of schema deletion on Workload-AUC.

by passing different weighted combinations of all input tokens to each expert. In addition to these works on MoE architectures and training strategies, recent years have witnessed the application of MoE in vision-related and multimodal predictive tasks [46]. However, our research delves into the potential of MoE in structured data analytics. We closely combine it with database data analytics, dynamically selecting necessary experts corresponding to the filter conditions in SQL queries.

In-Database Machine Learning involves executing machine learning within the database. MADlib [23] is an open-source library providing SQL-based ML functions in PostgreSQL. Google ML library[3], and Microsoft’s SQL Server Machine Learning Services [4] offer SQL APIs for ML functions on Oracle, bigquery, and Microsoft SQL Server, respectively. However, they only support feeding the static model with tuples in a data table and do not adapt feeding it with meta-information such as query encoding to customize the model for inference. Adding this support requires significant changes to their infrastructure, which is either non-trivial or unfeasible since infrastructure code is not accessible. Therefore, none of them are directly comparable. Complementary to LEADS, we propose an effective in-database model selection technique, TRAILS, in [55], and we will report their integration in the future.

7 CONCLUSIONS

In this paper, we propose a novel SQL-aware dynamic model slicing technique called LEADS. We enhance the general model with the Mixture of Experts (MoE) technique and devise a SQL-aware gating network to dynamically customize a sliced model given the propositional formula in the user’s SQL query. We further encapsulate LEADS into an in-database inference extension for PostgreSQL. In the implementation, we incorporate three key optimizations to accelerate the in-database inference process. Extensive experiments on five real-world datasets show that LEADS consistently outperforms four baseline models, and the in-database inference extension significantly reduces inference time compared to traditional approaches. We have integrated LEADS into NeurDB, our ongoing implementation of an AI-powered autonomous data system.

Acknowledgement:

This research work is supported by Singapore Ministry of Education Academic Research Fund Tier 3 under MOE’s official grant number MOE2017-T3-1-007. Yuncheng Wu’s work is supported by National Key Research and Development Program of China (Grant No.2023YFB4503600).

REFERENCES

- [1] 2000. Census-Income (KDD). UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5N30T>.
- [2] 2023. gprx:Framework for developing PostgreSQL extensions in Rust. <https://github.com/pgcentralfoundation/pgrx>. Accessed: October 10, 2023.
- [3] 2023. Machine learning models in BigQuery ML. <https://cloud.google.com/bigquery/docs/create-machine-learning-model>. Accessed: October 10, 2023.
- [4] 2023. Microsoft SQL MLs. <https://learn.microsoft.com/en-us/sql/machine-learning/sql-server-machine-learning-services?view=sql-server-2017>. Accessed: October 10, 2023.
- [5] Shun-ichi Amari. 1993. Backpropagation and stochastic gradient descent method. *Neurocomputing* 5, 4-5 (1993), 185–196.
- [6] KirillOdintsov Anna Montoya, inversion and Martin Kotek. 2018. Home Credit Default Risk. <https://kaggle.com/competitions/home-credit-default-risk>
- [7] Arvind Arasu and Hector Garcia-Molina. 2003. Extracting structured data from web pages. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*. 337–348.
- [8] Gökhan Bakır. 2007. *Predicting structured data*. MIT press.
- [9] Shaofeng Cai, Gang Chen, Beng Chin Ooi, and Jinyang Gao. 2019. Model Slicing for Supporting Complex Analytics with Elastic Inference Cost and Resource Constraints. *Proceedings of the VLDB Endowment* 13, 2 (2019), 86–99.
- [10] Shaofeng Cai, Kaiping Zheng, Gang Chen, HV Jagadish, Beng Chin Ooi, and Meihui Zhang. 2021. Arm-net: Adaptive relation modeling network for structured data. In *SIGMOD*. 207–220.
- [11] WeiYu Cheng, Yanyan Shen, and Linpeng Huang. 2020. Adaptive factorization network: Learning adaptive-order feature interactions. In *AAAI*, Vol. 34. 3609–3616.
- [12] John Clore and Beata. 2014. Diabetes 130-US Hospitals for Years 1999-2008. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5230J>.
- [13] Gonçalo M Correia, Vlad Niculae, and André FT Martins. 2019. Adaptively sparse transformers. *arXiv preprint arXiv:1909.00015* (2019).
- [14] Daniel Crankshaw, Xin Wang, Giulio Zhou, Michael J. Franklin, Joseph E. Gonzalez, and Ion Stoica. 2017. Clipper: A Low-Latency Online Prediction Serving System. In *NSDI*, Aditya Akella and Jon Howell (Eds.). 613–627.
- [15] Jesse Davis and Mark Goodrich. 2006. The relationship between Precision-Recall and ROC curves. In *ICML*. 233–240.
- [16] William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research* 23, 1 (2022), 5232–5270.
- [17] Mudasir A Ganaie, Minghui Hu, Ashwani Kumar Malik, Muhammad Tanveer, and Ponnuthurai N Suganthan. 2022. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence* 115 (2022), 105151.
- [18] Matt W Gardner and SR Dorling. 1998. Artificial neural networks (the multilayer perceptron)—a review of applications in the atmospheric sciences. *Atmospheric environment* 32, 14-15 (1998), 2627–2636.
- [19] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. 6.2. 2.3 softmax units for multinoulli output distributions. *Deep learning* 180 (2016).
- [20] Jim Gray, Surajit Chaudhuri, Adam Bosworth, Andrew Layman, Don Reichart, Murali Venkatrao, Frank Pellow, and Hamid Pirahesh. 1997. Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub-totals. *Data mining and knowledge discovery* 1 (1997), 29–53.
- [21] Jiaao He, Jiezhong Qiu, Aohan Zeng, Zhilin Yang, Jidong Zhai, and Jie Tang. 2021. Fastmoe: A fast mixture-of-expert training system. *arXiv preprint arXiv:2103.13262* (2021).
- [22] Xiangnan He and Tat-Seng Chua. 2017. Neural factorization machines for sparse predictive analytics. In *SIGIR*. 355–364.
- [23] Joe Hellerstein, Christopher Ré, Florian Schoppmann, Daisy Zhe Wang, Eugene Fratkin, Aleksander Gorajek, Kee Siong Ng, Caleb Welton, Xixuan Feng, Kun Li, et al. 2012. The MADlib analytics library or MAD skills, the SQL. *arXiv preprint arXiv:1208.4165* (2012).
- [24] Xin Huang, Ashish Khetan, Milan Cvitkovic, and Zohar Karnin. 2020. Tabtransformer: Tabular data modeling using contextual embeddings. *arXiv preprint arXiv:2012.06678* (2020).
- [25] Peng Jia, Shaofeng Cai, Beng Chin Ooi, Pinghui Wang, and Yiyuan Xiong. 2023. Robust and Transferable Log-based Anomaly Detection. *Proceedings of the ACM on Management of Data* 1, 1 (2023), 1–26.
- [26] Michael I Jordan and Robert A Jacobs. 1994. Hierarchical mixtures of experts and the EM algorithm. *Neural computation* 6, 2 (1994), 181–214.
- [27] Mahmoud Abo Khamis, Hung Q Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. 2020. Learning models over relational data using sparse tensors and functional dependencies. *TODS* 45, 2 (2020), 1–66.
- [28] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [29] Hyoukjun Kwon, Prasanth Chatarasi, Michael Pellauer, Angshuman Parashar, Vivek Sarkar, and Tushar Krishna. 2019. Understanding reuse, performance, and hardware cost of dnn dataflow: A data-centric approach. In *MICRO*. 754–768.
- [30] Dmitry Lepikhin, Hyoukjoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668* (2020).
- [31] Side Li, Lingjiao Chen, and Arun Kumar. 2019. Enabling and optimizing non-linear feature interactions in factorized linear algebra. In *SIGMOD*. 1571–1588.
- [32] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. 2018. xdeepfm: Combining explicit and implicit feature interactions for recommender systems. In *KDD*. 1754–1763.
- [33] Bin Lin, Zhenyu Tang, Yang Ye, Jiayi Cui, Bin Zhu, Peng Jin, Junwu Zhang, Munan Ning, and Li Yuan. 2024. MoE-LLaVA: Mixture of Experts for Large Vision-Language Models. *CoRR abs/2401.15947* (2024). <https://doi.org/10.48550/ARXIV.2401.15947>
- [34] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. 2017. Learning efficient convolutional networks through network slimming. In *ICCV*. 2736–2744.
- [35] Andre Martins and Ramon Astudillo. 2016. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *International conference on machine learning*. PMLR, 1614–1623.
- [36] Raghunath Othayoth Nambiar and Meikel Poess. 2006. The Making of TPC-DS. In *VLDB*, Vol. 6. 1049–1058.
- [37] Guillermo Navas-Palencia. 2020. Optimal binning: mathematical programming formulation. *arXiv preprint arXiv:2001.08025* (2020).
- [38] Milos Nikolic, Haozhe Zhang, Ahmet Kara, and Dan Olteanu. 2020. F-IVM: learning over fast-evolving relational data. In *SIGMOD*. 2773–2776.
- [39] Beng Chin Ooi, Shaofeng Cai, Gang Chen, Yanyan Shen, Kian-Lee Tan, Yuncheng Wu, Xiaokui Xiao, Naili Xing, Cong Yue, Lingze Zeng, and Meihui Zhang. 2024. NeurDB: an AI-powered autonomous data system. *Science China Information Sciences* 67, 10 (2024), 200901.
- [40] OpenAI. 2023. GPT-4 Technical Report. *CoRR abs/2303.08774* (2023). <https://doi.org/10.48550/ARXIV.2303.08774>
- [41] Kwanghyun Park, Karla Saur, Dalitso Banda, Rathijit Sen, Matteo Interlandi, and Konstantinos Karanasos. 2022. End-to-end Optimization of Machine Learning Prediction Queries. In *SIGMOD*, Zachary G. Ives, Angela Bonifati, and Amr El Abbadi (Eds.). ACM, 587–601.
- [42] Ben Peters, Vlad Niculae, and André FT Martins. 2019. Sparse Sequence-to-Sequence Models. In *ACL*. <https://www.aclweb.org/anthology/P19-1146>
- [43] Joan Puigcerver, Carlos Riquelme, Basil Mustafa, and Neil Houlsby. 2023. From sparse to soft mixtures of experts. *arXiv preprint arXiv:2308.00951* (2023).
- [44] S Brintha Rajakumari and C Nalini. 2014. An efficient data mining dataset preparation using aggregation in relational database. *Indian Journal of Science and Technology* 7 (2014), 44.
- [45] Steffen Rendle. 2013. Scaling factorization machines to relational data. *Proceedings of the VLDB Endowment* 6, 5 (2013), 337–348.
- [46] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. 2021. Scaling vision with sparse mixture of experts. *NeurIPS* (2021), 8583–8595.
- [47] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538* (2017).
- [48] Will Cukierski Steve Wang. 2014. Click-Through Rate Prediction. <https://kaggle.com/competitions/avazu-ctr-prediction>
- [49] John N Strack et al. 2014. Impact of HbA1c measurement on hospital readmission rates: analysis of 70,000 clinical database patient records. *BioMed research international* 2014 (2014).
- [50] Constantino Tsallis. 1988. Possible generalization of Boltzmann-Gibbs statistics. *Journal of statistical physics* 52 (1988), 479–487.
- [51] Johannes Von Oswald, Christian Henning, Benjamin F Grewe, and João Sacramento. 2019. Continual learning with hypernetworks. *arXiv preprint arXiv:1906.00695* (2019).
- [52] Jian-Wei Wang, Li-Li Rong, Qiu-Hong Deng, and Ji-Yong Zhang. 2010. Evolving hypernetwork model. *The European Physical Journal B* 77 (2010), 493–498.
- [53] Wei Wang, Jinyang Gao, Meihui Zhang, Sheng Wang, Gang Chen, Teck Khim Ng, Beng Chin Ooi, Jie Shao, and Moaz Reyad. 2018. Rafiki: Machine Learning as an Analytics Service System. *Proceedings of the VLDB Endowment* 12, 2 (2018), 128–140.
- [54] Steve Waterhouse, David MacKay, and Anthony Robinson. 1995. Bayesian methods for mixtures of experts. *NIPS* 8 (1995), 351–357.
- [55] Naili Xing, Shaofeng Cai, Gang Chen, Zhaojing Luo, Beng Chin Ooi, and Jian Pei. 2024. Database native model selection: Harnessing deep neural networks in database systems. *Proceedings of the VLDB Endowment* 17, 5 (2024), 1020–1033.
- [56] Lijie Xu, Shuang Qiu, Binhang Yuan, Jiawei Jiang, Cedric Renggli, Shaoduo Gan, Kaan Kara, Guoliang Li, Ji Liu, Wentao Wu, et al. 2022. In-Database Machine Learning with CorgiPile: Stochastic Gradient Descent without Full Data Shuffle. In *SIGMOD*. 1286–1300.
- [57] Lei Xu and Kalyan Veeramachaneni. 2018. Synthesizing tabular data using generative adversarial networks. *arXiv preprint arXiv:1811.11264* (2018).

- [58] I-Cheng Yeh. 2016. Default of credit card clients. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C55S3H>.
- [59] I-Cheng Yeh and Che-hui Lien. 2009. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert systems with applications* 36, 2 (2009), 2473–2480.
- [60] Yuhao Zhang, Frank Mcquillan, Nandish Jayaram, Nikhil Kak, Ekta Khanna, Orhan Kislal, Domino Valdano, and Arun Kumar. 2021. Distributed deep learning on data systems: a comparative analysis of approaches. *Proceedings of the VLDB Endowment* 14, 10 (2021).
- [61] Zhanhao Zhao, Shaofeng Cai, Haotian Gao, Hexiang Pan, Siqi Xiang, Naili Xing, Gang Chen, Beng Chin Ooi, Yanyan Shen, Yuncheng Wu, and Meihui Zhang. 2025. NeurDB: On the Design and Implementation of an AI-powered Autonomous Database. In *CIDR*.
- [62] Kaiping Zheng, Shaofeng Cai, Horng Ruey Chua, Wei Wang, Kee Yuan Ngiam, and Beng Chin Ooi. 2020. Tracer: A framework for facilitating accurate and interpretable analytics for high stakes applications. In *SIGMOD*. 1747–1763.
- [63] Yanqi Zhou, Tao Lei, Hanxiao Liu, Nan Du, Yanping Huang, Vincent Zhao, Andrew M Dai, Quoc V Le, James Laudon, et al. 2022. Mixture-of-experts with expert choice routing. *NeurIPS* (2022), 7103–7114.