



ADAPTIVE COST MODEL EVOLUTION THROUGH ONLINE LEARNING: A SELF-TUNING FRAMEWORK FOR DYNAMIC DATABASE OPTIMIZATION

Mohamed Goismi^{*1}, Moustafa Maaskri², Mohamed Debbab³ and Djamel Seghier⁴

^{1,3}Department of Science and Technology, University Ibn Khaldoun – Tiaret, Algeria.

²Department of Electrical Engineering, University Ibn Khaldoun – Tiaret, Algeria.

⁴Department of Computer Science, University Ibn Khaldoun – Tiaret, Algeria.

¹<http://orcid.org/0000-0001-9535-1142>, ²<http://orcid.org/0009-0000-5268-1133>,

³<http://orcid.org/0009-0003-9921-0824>, ⁴<http://orcid.org/0000-0003-1680-5033>

Email: ^{*}mohamed.goismi@univ-tiaret.dz, moustafa.maaskri@univ-tiaret.dz, mohamed.debbab@univ-tiaret.dz, djamal.seghier@univ-tiaret.dz

ARTICLE INFO

Article History

Recebido: 21 de Janeiro de 2026

Revisado: 22 de Março de 2026

Aceito: 22 de abril de 2026

Publicado: 29 de maio de 2026

Keywords:

Adaptive Cost Models, Online Learning, Reinforcement Learning, Query Optimization, Self-Tuning Databases, Meta-Learning, Database Performance.

ABSTRACT

Database cost models (CMs) are fundamental to query optimization, yet they suffer from accuracy degradation as system workloads and hardware configurations evolve. Traditional cost models rely on static calibration approaches that fail to adapt to dynamic operational environments. This paper introduces AdaptCM, a novel framework that leverages online learning and reinforcement learning techniques to enable continuous, autonomous cost model evolution. Unlike existing approaches that require manual recalibration or periodic offline updates, AdaptCM implements a self-tuning mechanism that learns from query execution feedback in realtime, automatically detecting model drift and adapting cost estimations without system disruption. We propose a multi-agent architecture where specialized learning agents manage different cost components (I/O, CPU, memory, network) and coordinate through a meta-learning controller. Our experimental evaluation on TPC-H, TPC-DS, and real-world production workloads demonstrates that AdaptCM achieves 23-41% higher prediction accuracy compared to PostgreSQL's native cost model and 15-28% improvement over state-of-the-art learned cost models. The framework reduces calibration overhead by 94% while maintaining model accuracy within 5% means relative error across diverse workload patterns and hardware configurations. Furthermore, AdaptCM's adaptation latency remains under 50ms, making it suitable for production deployment in high-throughput database systems.



Copyright ©2026 by authors and Galileo Institute of Technology and Education of the Amazon (ITEGAM). This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

I. INTRODUCTION

Query optimization in database management systems (DBMS) fundamentally depends on the accuracy of cost models that estimate the execution cost of alternative query plans [1], [2]. These cost models guide the query optimizer in selecting the most efficient execution strategy by predicting resource consumption—typically measured in terms of I/O operations, CPU cycles, memory usage, and network transfers. However, the effectiveness of traditional cost models has been increasingly challenged by three critical factors: (1) the dynamic nature of modern workloads characterized by shifting access patterns and data distributions [3], (2) the proliferation of heterogeneous hardware configurations including multi-core processors, NVMe storage, and specialized accelerators [4], [5], and (3) the complexity of cloud-native deployments where resource contention and virtualization introduce unpredictable performance variability [6].

1.1 MOTIVATION AND PROBLEM STATEMENT

Traditional database cost models employ analytical formulas derived from first principles or calibrated through offline benchmarking [7]. While this approach has served well in relatively static environments, it exhibits several fundamental limitations in contemporary database systems:

Static Calibration Inadequacy: Cost models are typically calibrated once during system installation or after major configuration changes. This static calibration fails to capture temporal variations in system behavior caused by workload evolution, data distribution shifts, or hardware aging effects [1]. Research indicates that cost model accuracy can degrade by 40-60% within months of initial calibration in production environments with dynamic workloads.

Manual Intervention Overhead: Maintaining cost model accuracy requires periodic manual recalibration—a laborintensive process demanding specialized expertise in database internals, profiling methodologies, and statistical analysis [8]. Enterprise DBAs report spending 15-25% of maintenance time on performance tuning activities directly or indirectly related to cost model calibration.

Coarse-Grained Adaptation: Existing approaches to cost model adaptation either apply global adjustments that overlook component-specific drift patterns or require complete model reconstruction, resulting in extended downtime and significant computational overhead [9].

Limited Feedback Integration: Current cost models rarely leverage runtime execution feedback systematically. While some systems collect execution statistics, this information is typically used for retrospective analysis rather than prospective model improvement [3]. Recent advances in learned cost models using deep learning have shown promise in improving initial accuracy [2], [10], [11]. However, these approaches inherit limitations from supervised learning paradigms: they require extensive offline training with representative workloads, struggle with distribution shift, and demand periodic retraining to maintain accuracy—essentially transferring the calibration burden from analytical to learned models without addressing the fundamental adaptation challenge.

1.2 KEY CONTRIBUTIONS

This paper presents AdaptCM, a novel framework for autonomous, continuous cost model evolution through online learning. Our approach fundamentally differs from existing solutions by treating cost model maintenance as an online learning problem rather than a periodic calibration task. The key contributions of this work include:

1. **Online Learning Architecture:** We introduce a multi-agent reinforcement learning architecture where specialized agents manage cost estimation for different system components (I/O, CPU, memory, network). These agents learn continuously from query execution feedback, enabling fine-grained adaptation to component-specific performance characteristics and drift patterns.
2. **Meta-Learning Controller:** We propose a metalearning controller that coordinates agent behaviors, detects model drift using statistical process control techniques, and dynamically adjusts learning rates to balance exploration-exploitation tradeoffs. This controller implements adaptive trust mechanisms that weight agent predictions based on recent accuracy history.
3. **Non-Disruptive Adaptation:** AdaptCM implements incremental model updates with dual-model validation, ensuring that adaptations improve rather than degrade performance. The framework maintains shadow models for validation and implements automatic rollback mechanisms when adaptations fail to meet quality thresholds.
4. **Workload-Aware Learning:** Our system incorporates workload characterization techniques that identify query patterns and adjust learning strategies accordingly. This enables the framework to handle mixed workloads (OLTP/OLAP) and rapid workload transitions without catastrophic forgetting.
5. **Comprehensive Evaluation:** We conduct extensive experiments on multiple benchmarks (TPC-H, TPC-DS, JOB) and real-world production traces from three enterprise environments, demonstrating 23-41% accuracy improvements over traditional models and 15-28% improvements over state-of-the-art learned models, while reducing calibration overhead by 94%.

1.3 PAPER ORGANIZATION

The remainder of this paper is organized as follows. Section II reviews related work in cost model design, learned query optimization, and online learning for databases. Section III presents the detailed architecture of AdaptCM, including the multi-agent learning framework, meta-learning controller, and drift detection mechanisms. Section IV describes the online learning algorithms and adaptation strategies. Section V presents our comprehensive experimental evaluation across multiple benchmarks and production workloads. Section VI discusses the implications, limitations, and future research directions. Finally, Section VII concludes the paper. It must define the problem and importance of the research carried out, it presents a (not very extensive) review of the literature about the article, including the authors' contributions to the state of the art. If you use abbreviations or acronyms, first write the words that identify them and then, in parentheses, the acronym. This set also establishes the research question, the objectives of the work and hypothesis, if necessary, the importance and limitations of the study. Establishes the method used at work. It is written in the present tense.

II. RELATED WORK AND BACKGROUND

II.1 TRADITIONAL COST MODELS

Cost Models Database cost models have evolved significantly since the early work on System R [12]. Traditional approaches rely on analytical formulas that model system components based on physical design principles. These models typically decompose query execution cost into components: I/O cost (disk accesses), CPU cost (tuple processing), and increasingly, memory and network costs for distributed systems [7]. The PostgreSQL cost model [13] represents a mature implementation of this approach, utilizing over 20 configurable parameters including seq page cost, random page cost, and cpu tuple cost. However, studies by [1] demonstrate that even carefully calibrated traditional models exhibit significant estimation errors, with Q-errors (ratio of estimated to actual cost) frequently exceeding 100× for complex queries. Commercial systems like Oracle and SQL Server have extended basic models with histogram-based cardinality estimation and adaptive query processing features [9]. Oracle's Adaptive Query Optimization monitors execution statistics and can adjust subsequent executions, but this adaptation operates at the query plan level rather than the cost model level.

II.2 LEARNED COST MODELS

The past five years have witnessed significant interest in applying machine learning to query optimization. In [2] pioneered the use of deep learning for cardinality estimation, demonstrating that neural networks can outperform traditional methods by learning from query execution history. This work inspired subsequent research into learned cost models. MSCN (Multi-Set Convolutional Networks) [14] introduced a specialized neural architecture for join cardinality estimation, achieving significant improvements over traditional methods on the JOB benchmark. However, the model requires extensive offline training and offers limited mechanisms for online adaptation. According to [15] proposed a bandit-based approach to query optimization that learns from execution feedback using Thompson sampling. While Bao addresses the online learning challenge, it focuses on plan selection rather than cost estimation, treating the cost model as a black box. In [10] introduced an end-to-end learning framework that jointly optimizes cardinality estimation and cost modeling. Despite impressive results, the approach requires periodic retraining with accumulated execution logs—a process that can take hours for large datasets and doesn't address real-time adaptation. Recent work by [16] provides a comprehensive benchmark comparing learned cardinality estimators, highlighting that no single approach dominates across all workloads and that models trained on one dataset often fail to generalize to others—a finding that motivates our focus on continuous adaptation.

II.3 ONLINE LEARNING AND ADAPTIVE SYSTEMS

Online learning algorithms update models incrementally as new data arrives, without requiring complete retraining [17]. Techniques such as stochastic gradient descent, online gradient descent, and adaptive learning rates (Adam, AdaGrad) form the foundation of many online learning systems. Reinforcement learning (RL) provides a framework for learning optimal policies through interaction with an environment [18]. Multi-armed bandits, a simplified RL setting, have been successfully applied to various online optimization problems. Contextual bandits extend this framework by incorporating context information, making them particularly suitable for adaptive systems where decisions depend on current state. In database systems, online learning has been applied primarily to workload management and resource allocation rather than cost modeling. CDBTune [19] uses deep reinforcement learning for database configuration tuning but operates on a coarser timescale (hours to days) than required for cost model adaptation.

II.4 SELF-TUNING DATABASE SYSTEMS

Self-tuning capabilities have been a research goal since the early 2000s. Microsoft's AutoAdmin project [20] pioneered automated physical design tuning, while IBM's LEO (Learning Optimizer) [9] introduced techniques for learning improved cardinality estimates from query feedback. More recent work includes adaptive indexing [21], which incrementally builds indexes based on workload patterns, and adaptive query processing techniques [22] that adjust execution strategies based on runtime observations. However, these approaches focus on physical design and execution strategies rather than fundamental cost model adaptation.

II.5 RESEARCH GAP

While existing work has made significant progress in both learned cost models and adaptive database systems, a critical gap remains: there exists no comprehensive framework for continuous, autonomous cost model evolution that addresses the full spectrum of challenges including component-specific adaptation, workload-aware learning, drift detection, and non-disruptive updates. AdaptCM addresses this gap by synthesizing techniques from online learning, reinforcement learning, and meta-learning into an integrated framework specifically designed for cost model adaptation in production database environments.

III. ADAPTCM ARCHITECTURE

III.1 SYSTEM OVERVIEW

AdaptCM implements a layered architecture consisting of four primary components (Figure 1): (1) Query Execution Monitor, (2) Multi-Agent Learning Framework, (3) Meta-Learning Controller, and (4) Cost Model Interface. This design separates concerns while enabling tight integration with existing database systems. The Query Execution Monitor captures execution statistics without introducing significant runtime overhead. For each executed query, it records: (1) estimated costs from the current model, (2) actual execution metrics (elapsed time, I/O operations, CPU cycles, memory usage), (3) query features (operators, cardinalities, selectivities), and (4) system state (buffer pool status, concurrent queries, resource utilization). This information flows to the learning framework through a lightweight, lock-free ring buffer implementation that minimizes performance impact. The Multi-Agent Learning Framework decomposes cost model adaptation into component-specific learning tasks. We employ four specialized agents: IOAgent, CPUAgent, MemoryAgent, and NetworkAgent. Each agent maintains a neural network model that predicts costs for its respective component and updates this model based on execution feedback. Agents implement modified Q-learning algorithms adapted for continuous action spaces (cost predictions) with experience replay for stability. The Meta-Learning Controller coordinates agent activities and manages the overall adaptation process. It implements: (1) drift detection using Page-Hinkley test and ADWIN (Adaptive Windowing), (2) learning rate adaptation based on prediction accuracy trends, (3) agent trust scoring that weights predictions according to recent performance, and (4) model validation using shadow models to ensure adaptations improve accuracy. The controller employs a meta-policy learned through meta-reinforcement learning that optimizes the trade-off between model stability and adaptation speed.




Figure 1: AdaptCM System Architecture showing the multi-agent learning framework, meta-learning controller, and integration with database query optimizer.

Source: Authors, (2026).

The Meta-Learning Controller coordinates agent activities and manages the overall adaptation process. It implements: (1) drift detection using Page-Hinkley test and ADWIN (Adaptive Windowing), (2) learning rate adaptation based on prediction accuracy trends, (3) agent trust scoring that weights predictions according to recent performance, and (4) model validation using shadow models to ensure adaptations improve accuracy. The controller employs a meta-policy learned through meta-reinforcement learning that optimizes the trade-off between model stability and adaptation speed. The Cost Model Interface provides seamless integration with database optimizers. It implements the standard cost model API expected by the query optimizer while routing cost estimation requests to the appropriate learning agents. The interface maintains both production and shadow models, enabling A/B testing of model updates before deployment. It also implements fallback mechanisms that revert to traditional cost formulas when agent confidence scores fall below thresholds.

III.2 MULTI-AGENT LEARNING FRAMEWORK

III.2.1 Agent Architecture

Each learning agent consists of three neural networks: a value network $V_\theta(s)$ that estimates the quality of the current cost model state, a policy network $\pi_\phi(a|s)$ that determines cost adjustment actions, and a prediction network $C_\psi(q, s)$ that generates actual cost estimates for queries. The prediction network takes query features q and system state s as input and outputs cost predictions. The prediction network architecture employs a hybrid design combining: (1) a set-based deep learning model for handling variable-length query plans (similar to MSCN [14]), (2) feature crosses capturing interactions between query characteristics and system state, and (3) residual connections that preserve the traditional cost model baseline. This architecture enables the network to learn corrections to analytical models rather than learning from scratch, significantly improving sample efficiency.

Formally, for query q with traditional cost estimate $\hat{c}_{\text{traditional}}(q)$, the agent's prediction is:

$$\hat{c}_{\text{agent}}(q, s) = \hat{c}_{\text{traditional}}(q) + f_\psi(q, s) \quad (1)$$

Where f_ψ represents the learned correction function. This formulation ensures that the learned model defaults to traditional estimates when training data is limited.

III.2.2 Learning Algorithm

Agents implement a modified Deep Q-Network (DQN) algorithm [23] adapted for cost estimation. The state space includes query features (operator types, table sizes, selectivities), historical prediction errors, and system metrics (resource utilization, concurrent workload). The action space consists of cost adjustment factors applied to predictions.

Algorithm 1 Agent Learning Update

```

1: Input: Mini-batch  $\mathcal{B} = \{(s_i, q_i, \hat{c}_i, c_i, s'_i)\}$ 
2: Initialize: Prediction error  $\mathcal{L}_p = 0$ , Q-value loss  $\mathcal{L}_q = 0$ 
3: for each  $(s, q, \hat{c}, c, s')$  in  $\mathcal{B}$  do
4:   // Update prediction network
5:    $\hat{c}_{new} \leftarrow C_\psi(q, s)$ 
6:    $\mathcal{L}_p \leftarrow \mathcal{L}_p + \text{Huber}(\hat{c}_{new}, c)$ 
7:   // Update Q-network
8:    $r \leftarrow -\text{MRE}(\hat{c}, c)$  {Reward}
9:    $Q_{target} \leftarrow r + \gamma \max_{a'} Q_{\theta'}(s', a')$ 
10:   $\mathcal{L}_q \leftarrow \mathcal{L}_q + (Q_\theta(s, a) - Q_{target})^2$ 
11: end for
12: Update  $\psi$  using  $\nabla_\psi \mathcal{L}_p$ 
13: Update  $\theta$  using  $\nabla_\theta \mathcal{L}_q$ 
14: Soft-update target network:  $\theta' \leftarrow \tau\theta + (1 - \tau)\theta'$ 

```

Each agent maintains an experience replay buffer storing tuples $(s_t, q_t, \hat{c}_{t+1}, c_{t+1})$ where s_t is the system state, q_t is the query, $\hat{c}_t$ is the predicted cost, c_t is the observed execution cost, and s_{t+1} is the resulting state. The agent samples mini-batches from this buffer to update its networks:

The reward function is designed to minimize relative error while penalizing large absolute deviations:

$$r(\hat{c}, c_{actual}) = -\alpha \cdot \text{MRE}(\hat{c}, c_{actual}) - \beta \cdot I[|\hat{c} - c_{actual}| > \Delta] \quad (2)$$

Where MRE is mean relative error, I is an indicator function, and Δ is a threshold for acceptable absolute deviation. Hyperparameters α and β control the trade-off between relative and absolute accuracy.

III.2.3 Agent Coordination

Agents must coordinate because cost components exhibit dependencies—for example, memory pressure affects I/O patterns. We implement a communication protocol where agents share state information through a shared memory space updated after each learning step. The total cost estimate combines individual agent predictions:

$$\hat{c}_{total}(q, s) = \sum_i w_i(s) \cdot \hat{c}_i(q, s) \quad (3)$$

Where $w_i(s)$ are trust scores assigned by the meta-controller (Section III.3), and $\hat{c}_i$ are individual agent predictions. This weighted combination allows the system to adapt dynamically to varying component importance across different workload phases.

III.3 META-LEARNING CONTROLLER

The meta-controller implements a hierarchical control policy that manages agent behaviors and orchestrates the adaptation process. Its design addresses three key challenges: detecting when adaptation is necessary, determining optimal learning parameters, and validating that adaptations improve performance.

III.3.1 Drift Detection

AdaptCM implements a two-tiered drift detection mechanism. The fast detector uses a sliding window approach with the Page-Hinkley test [24] to identify abrupt changes in error distributions. For each agent, we maintain a cumulative sum of standardized prediction errors:

$$S_t = \sum_{i=1}^t (\epsilon_i - \bar{\epsilon} - \delta) \quad (4)$$

Where ϵ_i is the prediction error at time i , $\bar{\epsilon}$ is the historical mean error, and δ is a sensitivity parameter. Drift is detected when $S_t - m_t > \lambda$ where $m_t = \min_{1 \leq k \leq t} S_k$ and λ is a threshold. The slow drift detector uses ADWIN (Adaptive Windowing) [25] to identify gradual drift over longer timescales. ADWIN maintains a variable-length window of recent prediction errors and detects changes in the error distribution by comparing statistics of sub-windows. This dual approach enables AdaptCM to respond quickly to sudden workload changes while avoiding false alarms from transient fluctuations.

III.3.2 Learning Rate Adaptation

The meta-controller dynamically adjusts agent learning rates based on detection signals and recent performance trends. When drift is detected, learning rates increase to accelerate adaptation; during stable periods, rates decrease to prevent overfitting to noise. The adaptation policy is itself learned through meta-reinforcement learning.

We formulate learning rate selection as a contextual bandit problem where the context includes: (1) drift detection signals, (2) prediction error trends, (3) workload characteristics, and (4) agent confidence scores. The meta-controller maintains distribution over potential learning rates for each agent and selects rates using Thompson sampling:

$$\eta_i^{(t)} \sim P(\eta \mid \text{context}_t, \text{history}_{0:t-1}) \quad (5)$$

This probabilistic selection enables exploration of learning rate configurations while exploiting known good settings.

III.3.3 Trust Scoring and Model Selection

The meta-controller assigns trust scores to each agent based on recent prediction accuracy. Trust scores exponentially weight recent performance:

$$\text{trust}_i^{(t)} = \frac{\sum_{k=t-W}^t \gamma^{t-k} \cdot \text{accuracy}_i^{(k)}}{\sum_{k=t-W}^t \gamma^{t-k}} \quad (6)$$

where W is the window size and $\gamma \in (0, 1)$ is a decay factor. These scores determine agent influence in the final cost estimate (Equation 3) and trigger re-calibration when scores fall below thresholds.

III.4 NON-DISRUPTIVE ADAPTATION PROTOCOL

AdaptCM implements a multi-stage validation protocol to ensure that model updates improve rather than degrade performance:

Stage 1: Shadow Model Training. When adaptation is triggered, agents train shadow models using recent execution feedback while production models continue serving cost estimates.

Stage 2: Validation Testing. Shadow models are validated on a held-out set of recent queries. If shadow model accuracy exceeds production model accuracy by a threshold ϵ , proceed to Stage 3; otherwise, discard shadow models.

Stage 3: Gradual Rollout. Shadow models are promoted to production gradually using a canary deployment approach. Initially, 10% of cost estimation requests route to shadow models. If error rates remain acceptable, the percentage increases incrementally (10%, 25%, 50%, 75%, 100%) over subsequent time periods.

Stage 4: Monitoring and Rollback. Throughout rollout, the meta-controller monitors prediction errors. If errors exceed thresholds or drift detection fires, the system automatically rolls back to the previous production model.

This protocol ensures that adaptations are validated before full deployment and can be reverted quickly if problems arise—critical requirements for production database systems where cost model errors directly impact query performance.

IV. ONLINE LEARNING ALGORITHMS

IV.1 QUERY FEATURE EXTRACTION

Effective online learning requires informative feature representations that capture query characteristics relevant to execution cost. AdaptCM extracts features at three levels:

Operator-Level Features: For each physical operator in the query plan (e.g., sequential scan, index scan, hash join), we extract: operator type (one-hot encoded), input cardinality, output cardinality, selectivity, and operator-specific parameters (e.g., hash table size, sort buffer size). These features are aggregated using a set-based architecture similar to Deep Sets [26] to handle variable-length plans.

Plan-Level Features: Global plan characteristics including total number of operators, plan tree depth, number of joins, maximum table size, estimated intermediate result sizes, and parallelism degree. We also include a learned embedding of the plan structure using a Graph Neural Network that captures operator dependencies.

Context Features: System state information at query submission time: buffer pool hit rate, number of concurrent queries, CPU utilization, I/O queue depth, and memory pressure indicators. Context features enable the model to account for resource contention effects that traditional models often ignore.

The complete feature vector x_q for query q concatenates these three levels. To handle high dimensionality while enabling efficient online updates, we apply random projections to reduce feature space dimensions from 500 to 100 while preserving approximate distances (Johnson-Lindenstrauss lemma).

IV.2 INCREMENTAL MODEL UPDATES

Agent prediction networks are updated incrementally using streaming variants of gradient descent. For each executed query, we compute the prediction error and perform a single gradient update:

$$\psi_{t+1} = \psi_t - \eta_t \cdot \nabla_{\psi} \mathcal{L}(C_{\psi}(q_t, s_t), c_{\text{actual},t}) \quad (7)$$

Where $\mathcal{L}$ is the Huber loss function (robust to outliers) and η_t is the learning rate determined by the metacontroller. To prevent catastrophic forgetting - where the model loses accuracy on older workload patterns - we implement two complementary techniques:

Experience Replay with Prioritized Sampling: Each agent maintains a bounded replay buffer storing recent execution records. Updates sample mini-batches from this buffer, mixing recent experiences with historical ones. Sampling probabilities are proportional to prediction errors, ensuring the model focuses on difficult examples:

$$p(i) \propto |\epsilon_i|^\alpha + \epsilon_{\min} \quad (8)$$

Where α controls prioritization strength and $\epsilon_{\min}$ ensures all samples have non-zero probability. Elastic Weight Consolidation (EWC): Inspired by continual learning research [27], we add a regularization term that penalizes changes to parameters important for previous tasks:

$$\mathcal{L}_{\text{EWC}} = \mathcal{L}_{\text{current}} + \frac{\lambda}{2} \sum_i F_i(\psi_i - \psi_i^*)^2 \quad (9)$$

Where F_i is the Fisher information matrix diagonal (approximating parameter importance), ψ^* represents parameters after previous training phases, and λ controls regularization strength. This technique significantly improves model stability during workload transitions.

IV.3 HANDLING DATA DISTRIBUTION SHIFT

Production database workloads exhibit various forms of distribution shift: sudden changes (e.g., batch job execution), gradual drift (e.g., data growth), and periodic patterns (e.g., daily cycles). AdaptCM addresses these challenges through workload-aware adaptation strategies.

Workload Characterization: The system maintains a workload classifier that categorizes queries into types (OLTP, OLAP, mixed) based on features such as query complexity, data access patterns, and resource requirements. This classifier is implemented as a lightweight online clustering algorithm (Mini-Batch K-Means) that incrementally updates cluster centers as new queries arrive. For each workload cluster, AdaptCM maintains separate experience replay buffers and adaptation parameters. When workload distribution shifts from one cluster to another, the meta-controller adjusts learning strategies accordingly:

Periodic Pattern Recognition: Many workloads exhibit daily or weekly patterns (e.g., increased load during business hours). AdaptCM implements a time-aware extension where models learn time-varying correction factors. We augment query features with temporal encodings (hour of day, day of week) and train separate correction networks for identified periodic patterns.

Algorithm 2 Workload-Aware Adaptation

```

1: Input: Query  $q$ , current workload distribution  $\mathcal{W}_t$ 
2: Classify query:  $cluster \leftarrow \text{WorkloadClassifier}(q)$ 
3: Update distribution:  $\mathcal{W}_t[cluster] \leftarrow \mathcal{W}_t[cluster] + 1$ 
4: if  $\text{KL}(\mathcal{W}_t, \mathcal{W}_{t-1}) > \tau_{\text{shift}}$  then
5:   // Significant workload shift detected
6:   Increase learning rates:  $\eta_i \leftarrow \eta_i \cdot \alpha_{\text{shift}}$ 
7:   Adjust replay buffer sampling: favor  $cluster$  samples
8:   Trigger meta-controller reoptimization
9: end if
10: Route query to appropriate agent configuration
    
```

IV.4 MULTI-OBJECTIVE OPTIMIZATION

Cost model adaptation must balance multiple objectives: prediction accuracy, model stability, adaptation speed, and computational overhead. We formulate this as a multi-objective optimization problem and use Pareto optimization techniques to find suitable trade-offs. The meta-controller's policy network is trained to maximize a composite reward:

$$R_{\text{meta}} = w_1 \cdot R_{\text{accuracy}} + w_2 \cdot R_{\text{stability}} + w_3 \cdot R_{\text{efficiency}} \quad (10)$$

Where:

- $R_{\text{accuracy}} = -\text{MRE}$ (negative mean relative error)
- $R_{\text{stability}} = -\text{Var}(\epsilon_t)$ (negative error variance)
- $R_{\text{efficiency}} = -\frac{t_{\text{update}}}{t_{\text{budget}}}$ (ratio of update time to budget)

Weights w_1, w_2, w_3 are configurable to reflect deployment priorities. For production systems prioritizing stability, w_2 is increased; for systems requiring rapid adaptation, w_1 and w_3 dominate.

V. EXPERIMENTAL EVALUATION

V.1 EXPERIMENTAL SETUP

V.1.1 Systems and Configurations

We evaluated AdaptCM on PostgreSQL 14.5, comparing against four baselines:

1. **PostgreSQL Default:** Native cost model with default parameters;
2. **PostgreSQL Tuned:** Native model with parameters carefully calibrated using pg bench and manual tuning (requiring 16 hours of expert time);
3. **Neo [2]:** State-of-the-art learned cardinality estimator integrated with PostgreSQL’s cost model;
4. **Bao [15]:** Bandit-based query optimizer that learns from execution feedback.

Table 1: Benchmark Characteristics.

Benchmark	Scale Factor	Tables	Queries
TPC-H	10 GB	8	22
TPC-H	100 GB	8	22
TPC-DS	10 GB	24	99
JOB	3.8 GB	21	113
Production-A	45 GB	127	2,847
Production-B	218 GB	89	1,653
Production-C	512 GB	203	4,291

Source: Authors, (2026).

All experiments ran on identical hardware: Intel Xeon E5-2690 v4 (2.6 GHz, 14 cores), 128 GB DDR4 RAM, 1 TB NVMe SSD, running Ubuntu 20.04 LTS. The database buffer pool was configured to 32 GB, with remaining PostgreSQL parameters at defaults unless otherwise noted. AdaptCM was implemented as a PostgreSQL extension using: PyTorch 1.12 for neural networks, PostgreSQL hooks for execution monitoring, and shared memory segments for agent communication. The prediction network for each agent consists of 4 hidden layers (256, 128, 64, 32 units) with ReLU activations. The meta-controller policy network has 3 hidden layers (128, 64, 32 units). All networks were initialized using Kaiming initialization and trained with Adam optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.999$).

V.1.2 Benchmarks and Workloads

We evaluated on six workloads (Table 1):

TPC-H [28]: Standard decision support benchmark with star schema and analytical queries. We used scale factors 10 GB and 100 GB to evaluate scaling behavior.

TPC-DS [29]: More complex decision support benchmark with snowflake schema, featuring 99 queries with multiple joins, aggregations, and subqueries.

JOB (Join Order Benchmark) [1]: Real-world workload based on IMDB data with highly skewed distributions and complex join queries—specifically designed to challenge query optimizers.

Production Workloads: Three anonymized production database workloads from: (A) e-commerce platform (mixed OLTP/OLAP), (B) financial services application (primarily OLAP), and (C) SaaS analytics platform (time-series heavy). These workloads provided realistic distribution shifts and complex access patterns.

For each benchmark, we executed queries in random order to simulate realistic workload patterns. Production workloads were replayed using timestamps from original executions to preserve temporal patterns.

V.2 METRICS

We evaluated systems using established cost model quality metrics:

Mean Relative Error (MRE): Primary accuracy metric defined as:

$$\text{MRE} = \frac{1}{N} \sum_{i=1}^N \frac{|\hat{c}_i - c_i|}{c_i} \quad (11)$$

Q-Error: Multiplicative error metric capturing both over- and under-estimation:

$$\text{Q-Error}(q) = \max\left(\frac{\hat{c}}{c}, \frac{c}{\hat{c}}\right) \quad (12)$$

We report median and 95th percentile Q-errors.

Plan Quality: Percentage of queries where optimizer selected optimal or near-optimal plan (within 20% of optimal cost).

Adaptation Metrics: Time to achieve target accuracy after workload shift, overhead of learning updates, and memory footprint.

V.3 OVERALL ACCURACY RESULTS

Table 2 presents overall accuracy across all benchmarks. AdaptCM achieved lowest MRE on 6 of 7 workloads, with improvements of 23-41% over PostgreSQL default, 15-28% over PostgreSQL tuned, 18-35% over Neo, and 15-28% over Bao. Figure 2 shows Q-error distributions for TPC-DS. AdaptCM significantly reduced both median and tail Qerrors. The 95th percentile Q-error for AdaptCM was 8.7× compared to 24.3× for PostgreSQL tuned, 16.2× for Neo, and 12.4× for Bao—demonstrating more consistent accuracy across diverse query types. The performance advantage of AdaptCM is particularly pronounced on production workloads where data distributions and access patterns evolve over time. On Production-C (the largest and most dynamic workload),




Figure 2: Q-Error distributions on TPC-DS workload. AdaptCM achieves lower median and significantly better tail performance (95th percentile) compared to baselines.

Source: Authors, (2026).

Table 2: Mean Relative Error (%) - Lower is Better.

Workload	PG Def.	PG Tuned	Neo	Bao	AdaptCM
TPC-H 10 GB	34.7	28.3	25.1	22.9	18.2
TPC-H 100 GB	42.3	35.7	31.4	28.6	21.7
TPC-DS 10 GB	51.8	43.2	38.7	35.4	26.8
JOB	67.4	54.9	49.3	45.7	39.8
Production-A	38.9	31.5	28.4	26.1	19.7
Production-B	45.2	37.8	33.9	30.7	23.4
Production-C	53.7	44.6	40.2	36.8	27.9
Average	47.7	39.4	35.3	32.3	25.4

Source: Authors, (2026).




Figure 3: Temporal evolution of MRE on Production-C workload over 30 days. AdaptCM maintains stable accuracy while baselines degrade as workload evolves.

Source: Authors, (2026).

AdaptCM maintained MRE below 30% throughout a 30-day evaluation period, while PostgreSQL tuned degraded from 35% to 52% MRE, Neo degraded from 32% to 47%, and Bao stabilized around 38% (Figure 3).

V.4 ADAPTATION BEHAVIOR

To evaluate AdaptCM's adaptation capabilities, we conducted controlled experiments introducing abrupt and gradual workload shifts.

V.4.1 Abrupt Workload Shifts

We transitioned from TPC-H queries to TPC-DS queries at hour 12 of a 24-hour run (Figure 4). AdaptCM detected the shift within 3.2 minutes (using Page-Hinkley test) and adapted to achieve stable accuracy (MRE < 25%) within 47 minutes. In contrast, Neo required complete retraining (taking 8.3 hours with accumulated execution logs), while Bao adapted more slowly, achieving stable accuracy after 2.4 hours. The rapid adaptation of AdaptCM results from its multi-agent architecture and meta-learning controller. When drift is detected, the controller increases learning rates for affected agents (primarily IOAgent and CPUAgent for this transition) while maintaining stability of others. Experience replays buffers prioritize recent examples during adaptation phase, accelerating convergence.

V.4.2 Gradual Drift

We simulated gradual data growth by adding 10% additional rows to TPC-H tables every 6 hours. AdaptCM's ADWIN detector identified the drift trend after the second growth period and gradually increased learning rates. MRE remained below 22% throughout the experiment (Figure 5), while PostgreSQL tuned accuracy degraded from 28% to 41% MRE as calibrated parameters became stale.

V.5 COMPONENT ANALYSIS

We conducted ablation studies to evaluate contributions of key AdaptCM components (Table 3). Multi-Agent Architecture: Decomposing cost estimation into component-specific agents provided significant benefits. The single-agent baseline (modelling all cost components with one network) achieved only 38.6% MRE—much worse than the full system's 26.8%. Component-specific agents learn more effectively because they focus on narrower prediction tasks and can specialize to component-specific characteristics.

Meta-Learning Controller: Removing the meta-controller and using fixed learning rates increased MRE to 32.4%, demonstrating the value of adaptive learning rate control and agent coordination.

Drift Detection: Without drift detection, the system cannot identify when to accelerate adaptation. This degraded MRE to 34.7%, particularly hurting performance during workload transitions where the system continued with slow learning rates inappropriate for rapid changes.

Experience Replay and EWC: Both techniques address catastrophic forgetting. Disabling experience replay increased MRE to 35.9%, while disabling EWC increased it to 31.2%. Combining both (in the full system) provides complementary benefits: experience replay enables learning from historical data, while EWC preserves important parameter configurations.

Workload Classification: The workload-aware adaptation strategy reduced MRE by 7.0 percentage points. Workload classification enables the system to maintain specialized models for different query types and adapt strategies when workload composition changes.




Figure 4: Response to abrupt workload shift from TPC-H to TPC-DS. AdaptCM detects shift quickly (vertical dashed line) and adapts within 47 minutes.

Source: Authors, (2026).




Figure 5: Handling gradual data growth (10% every 6 hours). AdaptCM maintains stable accuracy while PostgreSQL tuned degrades as parameters become stale.

Source: Authors, (2026).

Table 3: Ablation Study: Component Contributions (MRE % on TPC-DS).

Configuration	MRE (%)	Δ vs Full
AdaptCM (Full System)	26.8	-
Without Meta-Controller	32.4	+5.6
Without Drift Detection	34.7	+7.9
Without Experience Replay	35.9	+9.1
Without EWC Regularization	31.2	+4.4
Without Workload Classifier	33.8	+7.0
Single Agent (no decomposition)	38.6	+11.8
No Trust Scoring	30.1	+3.3
All Components Disabled	52.3	+25.5

Source: Authors, (2026).

V.6 PLAN QUALITY IMPACT

Improved cost model accuracy translates to better query plan selection. We evaluated plan quality by comparing optimizer choices against optimal plans determined through exhaustive search (feasible for JOB's smaller queries) or runtime measurement of top-k plans. Table 4 shows the percentage of queries where each system selected optimal or near-optimal plans (within 20% of optimal execution time). AdaptCM achieved optimal/near-optimal plan selection for 78.4% of JOB queries compared to 52.7% for PostgreSQL tuned and 64.3% for Bao. More importantly, AdaptCM significantly reduced catastrophically bad plan selections. For JOB queries, only 3.5% of AdaptCM's plans were more than $2\times$ slower than optimal, compared to 18.6% for PostgreSQL tuned, 12.4% for Neo, and 8.8% for Bao. This demonstrates that improved cost model accuracy provides the greatest benefit for complex queries where plan selection is most critical.

Table 4: Plan Quality: Percentage of Queries with.

Workload	PG Tuned	Neo	Bao	AdaptCM
TPC-H 10GB	68.2%	72.7%	77.3%	86.4%
TPC-DS 10GB	54.5%	59.6%	65.7%	73.7%
JOB	52.7%	57.5%	64.3%	78.4%
Average	58.5%	63.3%	69.1%	79.5%

Source: Authors, (2026).

V.7 OVERHEAD ANALYSIS

Production deployment requires that adaptation overhead remain acceptable. We measured three overhead components:

Query Execution Monitoring: Feature extraction and execution metric collection added median 1.2% overhead to query execution time (95th percentile: 3.4%). This overhead is dominated by plan tree traversal for feature extraction and remains constant regardless of query complexity.

Model Inference: Cost prediction through agent networks added median 0.8 ms per query (95th percentile: 2.3 ms). For queries with execution times exceeding 100 ms (the majority in analytical workloads), this represents !1% overhead. Even for fast OLTP queries (10-50 ms), overhead remained under 5%.

Model Updates: Learning updates occurred asynchronously after query execution, so they did not impact query latency. Update processing consumed average 3.2% of CPU time during periods of active adaptation, decreasing to !1% during stable periods. Peak memory usage for all agent networks, replay buffers, and meta-controller totaled 2.8 GB—acceptable for modern database servers. Table 5 compares overhead metrics across systems.

Table 5: System Overhead Comparison.

System	Query Overhead	CPU Usage	Memory (GB)
PostgreSQL Default	0%	0%	0
Neo	2.1%	0.8%	1.2
Bao	1.4%	1.2%	0.6
AdaptCM	1.9%	3.2%	2.8

Source: Authors, (2026).

V.8 CALIBRATION EFFORT REDUCTION

A key advantage of AdaptCM is eliminating manual calibration. We quantified calibration effort by measuring time required to achieve target accuracy:

PostgreSQL Tuned: Manual calibration required 16.3 hours of expert DBA time: profiling workload (4.5h), running calibration benchmarks (6.2h), parameter tuning (4.8h), and validation testing (0.8h). This process must repeat when workload or hardware changes.

Neo: Initial training required 12.7 hours: collecting execution traces (8.2h) and training neural networks (4.5h). Retraining after workload shift required 8.3 hours.

Bao: Bootstrap period to collect sufficient execution feedback required 6.4 hours of production traffic.

AdaptCM: Achieved target accuracy (MRE < 30%) within 2.1 hours using production traffic, with no manual intervention required. This represents a 94% reduction in calibration overhead compared to PostgreSQL tuned (Table 6).

Table 6: Calibration Time to Target Accuracy (MRE < 30%).

System	Time (hours)	Manual Effort
PostgreSQL Tuned	16.3	High
Neo	12.7	Medium
Bao	6.4	Low
AdaptCM	2.1	None

Source: Authors, (2026).

V.9 CROSS-SYSTEM PORTABILITY

We evaluated AdaptCM’s portability by training on PostgreSQL and testing on MySQL 8.0 (both using identical TPC-H 10GB data). While direct knowledge transfer between different database engines is challenging due to different internal implementations, AdaptCM’s architecture—which learns corrections to analytical models rather than replacing them—provides some transferability. After fine-tuning for 1.5 hours on MySQL execution traces, AdaptCM achieved 31.2% MRE on MySQL, compared to 36.8% for MySQL’s default cost model and 33.4% for manually tuned parameters. This demonstrates that learned corrections partially transfer across systems, though full accuracy requires system-specific adaptation.

VI. DISCUSSION

VI.1 KEY FINDINGS

Our evaluation demonstrates that online learning enables continuous, autonomous cost model evolution that significantly outperforms both traditional and learned static approaches. Several key findings emerge:

Multi-Agent Decomposition is Essential: Component-specific agents learn more effectively than monolithic models. Our ablation study showed that single-agent modelling increased MRE by 11.8 percentage points. Decomposition provides specialization benefits and enables fine-grained adaptation where different components drift at different rates.

Meta-Learning Enables Robust Adaptation: The meta-controller’s adaptive learning rate control and drift detection are critical for balancing exploration/exploitation. Fixed learning rates either adapt too slowly during workload shifts or overfit to noise during stable periods. Dynamic adjustment based on drift signals improved MRE by 5.6 percentage points. Workload Awareness Prevents Catastrophic

Forgetting: Maintaining workload-specific experience buffers and adaptation strategies reduced MRE by 7.0 percentage points. without workload classification, the model struggles during mixed-workload scenarios (e.g., Production-A with both OLTP and OLAP traffic).

Non-Disruptive Validation is Crucial: Shadow model validation and gradual rollout prevented 12 cases (across all experiments) where model updates would have degraded accuracy. Automatic rollback triggered 5 times, preventing performance regressions. Production deployment requires these safety mechanisms.

VI.2 LIMITATIONS AND FUTURE WORK

While AdaptCM demonstrates significant advantages, several limitations warrant discussion:

Cold Start Problem: AdaptCM requires a warm-up period (2-3 hours with typical workload) to collect sufficient execution feedback. During this period, accuracy may be inferior to carefully calibrated traditional models. Future work should investigate transfer learning from similar workloads to reduce cold start time.

Complex Query Coverage: Very complex queries with many joins and subqueries remain challenging. On JOB's most complex queries (.10 joins), MRE exceeded 40% for all systems including AdaptCM. Improved query representation techniques (e.g., transformer-based plan encoders) may help.

Distributed Systems: Our evaluation focused on single-node systems. Extending AdaptCM to distributed databases introduces additional challenges: network cost modeling, data locality effects, and distributed agent coordination. Preliminary experiments on Greenplum (PostgreSQL-based MPP database) show promise, but comprehensive evaluation is needed.

Adversarial Robustness: Online learning systems can be vulnerable to adversarial manipulation through carefully crafted query sequences. While unlikely in typical database environments, security-critical applications may require additional safeguards.

VI.3 DEPLOYMENT CONSIDERATIONS

Based on production trials at two partner organizations, we offer deployment recommendations:

Gradual Integration: Start by deploying AdaptCM in shadow mode (collecting feedback without affecting optimization) to build confidence before full deployment. This also populates experience replay buffers.

Monitoring and Alerting: Configure alerts for drift detection events, adaptation failures, and accuracy degradation. Integrate with existing database monitoring tools to provide visibility into adaptation behavior.

Fallback Configuration: Always maintain the ability to revert to traditional cost models. Configure confidence thresholds that trigger automatic fallback when prediction uncertainty exceeds limits.

Resource Provisioning: Allocate sufficient memory (3-4 GB) and CPU headroom (5-10% additional capacity) for adaptation overhead. Consider deploying learning components on separate nodes in multi-node configurations.

VII. CONCLUSION

This paper presented AdaptCM, a novel framework for continuous, autonomous database cost model evolution through online learning. By treating cost model maintenance as an online learning problem and employing a multi-agent architecture with meta-learning control, AdaptCM achieves 23-41% higher prediction accuracy compared to traditional approaches and 15-28% improvement over state-of-the-art learned models, while reducing calibration overhead by 94%. The key innovation lies in decomposing cost estimation into component-specific learning tasks managed by specialized agents, coordinated by a meta-learning controller that detects drift, adapts learning parameters, and validates model updates. This architecture enables AdaptCM to adapt continuously to evolving workloads and hardware configurations without manual intervention, addressing a fundamental limitation of existing cost modeling approaches. Our comprehensive evaluation on standard benchmarks and real-world production workloads demonstrates that AdaptCM maintains stable accuracy across diverse scenarios including abrupt workload shifts, gradual data drift, and mixed OLTP/OLAP environments. The framework's non-disruptive adaptation protocol ensures that model updates improve rather than degrade performance—a critical requirement for production deployment. Looking forward, several exciting research directions emerge: extending AdaptCM to distributed database systems, incorporating transfer learning to accelerate cold start, exploring attention-based query representations for improved complex query modeling, and investigating federated learning approaches that enable multiple database instances to share adaptation knowledge while preserving privacy. As database workloads continue to grow in complexity and diversity, adaptive cost models that learn continuously from production traffic represent a promising path toward truly self-tuning database systems. AdaptCM demonstrates that this vision is achievable with existing machine learning techniques, opening opportunities for next-generation database optimizers that maintain accuracy automatically in dynamic operational environments.

VIII. AUTHOR'S CONTRIBUTION

Conceptualization: Mohamed Goismi, Moustafa Maaskri, Mohamed Debbab and Djamel Seghier.

Methodology: Mohamed Goismi, Moustafa Maaskri, Mohamed Debbab and Djamel Seghier.

Investigation: Mohamed Goismi, Moustafa Maaskri, Mohamed Debbab and Djamel Seghier.

Discussion of results: Mohamed Goismi, Moustafa Maaskri, Mohamed Debbab and Djamel Seghier.

Writing – Original Draft: Mohamed Goismi, Moustafa Maaskri, Mohamed Debbab and Djamel Seghier.

Writing – Review and Editing: Mohamed Goismi, Moustafa Maaskri, Mohamed Debbab and Djamel Seghier.

Resources: Mohamed Goismi, Moustafa Maaskri, Mohamed Debbab and Djamel Seghier.

Supervision: Mohamed Goismi, Moustafa Maaskri, Mohamed Debbab and Djamel Seghier.

Approval of the final text: Mohamed Goismi, Moustafa Maaskri, Mohamed Debbab and Djamel Seghier.

IX. REFERENCES

- [1] V. Leis, A. Gubichev, A. Mirchev, P. Boncz, A. Kemper, and T. Neumann, "How good are query optimizers, really?" Proc. VLDB Endowment, vol. 9, no. 3, pp. 204–215, Nov. 2015. DOI:10.14778/2850583.2850594.
- [2] R. Marcus, P. Negi, H. Mao, C. Zhang, M. Alizadeh, T. Kraska, O. Papaemmanouil, and N. Tatbul, "Neo: A learned query optimizer," Proc. VLDB Endowment, vol. 12, no. 11, pp. 1705–1718, July 2019. DOI:10.14778/3342263.3342644.

- [3] W. Wu, Y. Chi, H. Hac'ig'um'u, s, and J. F. Naughton, "Towards predicting query execution time for concurrent and dynamic database workloads," Proc. VLDB Endowment, vol. 6, no. 10, pp. 925–936, Aug. 2013. DOI: 10.14778/2536206.2536219.
- [4] T. Kissinger, D. Habich, and W. Lehner, "Adaptive Energy-Control for In-Memory Database Systems," in Proc. ACM SIGMOD, 2018, pp. 351–364. DOI: 10.1145/3183713.3183756.
- [5] H. Nicholson, A. Raza, P. Chrysogelos, and A. Ailamaki, "HetCache: Synergising NVMe Storage and GPU acceleration for Memory-Efficient Analytics," in Proc. 13th Annu. Conf. Innovative Data Systems Research (CIDR '23), Amsterdam, The Netherlands, Jan. 2023.
- [6] X. Zhang, H. Wu, Z. Chang, S. Jian, Y. Xue, Y. Chen, J. Zhou, and L. Chen, "Facilitating database tuning with hyper-parameter optimization: A comprehensive experimental evaluation," Proc. VLDB Endowment, vol. 15, no. 9, pp. 1808–1821, May 2022. DOI: 10.14778/3538598.3538604.
- [7] S. Chaudhuri, "An overview of query optimization in relational systems," in Proc. 17th ACM SIGACT-SIGMOD-SIGART Symp. Principles of Database Systems, 1998, pp. 34–43. DOI:10.1145/275487.275492.
- [8] M. V. M. Lakshmi and S. Zhou, "Extracting meaningful data from damaged database storage," ACM Trans. Database Syst., vol. 30, no. 1, pp. 213–246, Mar. 2005. DOI: 10.1145/1061318.1061322.
- [9] M. Stillger, G. M. Lohman, V. Markl, and M. Kandil, "LEO - DB2's LEarning Optimizer," in Proc. 27th Int. Conf. Very Large Data Bases, 2001, pp. 19–28. DOI: 10.5555/645927.672244.
- [10] J. Sun and G. Li, "An end-to-end learningbased cost estimator," Proc. VLDB Endowment, vol. 13, no. 3, pp. 307–319, Nov. 2019. DOI:10.14778/3368289.3368296
- [11] J. Ortiz, M. Balazinska, J. Gehrke, and S. S. Keerthi, "Learning state representations for query optimization with deep reinforcement learning," in Proc. 2nd Workshop on Data Management for End-To-End Machine Learning, 2018, pp. 1–4. DOI: 10.1145/3209889.3209890.
- [12] P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price, "Access path selection in a relational database management system," in Proc. ACM SIGMOD Int. Conf. Management of Data, 1979, pp. 23–34. DOI: 10.1145/582095.582099.
- [13] PostgreSQL Global Development Group, "PostgreSQL 14.5 Documentation: Cost-based Vacuum Delay," 2023. [Online]. Available: <https://www.postgresql.org/docs/14/>
- [14] A. Kipf, T. Kipf, B. Radke, V. Leis, P. Boncz, and A. Kemper, "Learned cardinalities: Estimating correlated joins with deep learning," in Proc. 9th Biennial Conf. Innovative Data Systems Research, 2019. DOI: 10.48550/arXiv.1809.00677.
- [15] R. Marcus and O. Papaemmanouil, "Bao: Making learned query optimization practical," in Proc. ACM SIGMOD Int. Conf. Management of Data, 2021, pp. 1275–1288. DOI: 10.1145/3448016.3452838.
- [16] L. Woltmann, C. Hartmann, M. Thiele, D. Habich, and W. Lehner, "Cardinality estimation with local deep learning models," in Proc. 24th Int. Conf. Extending Database Technology, 2021, pp. 577–588. DOI: 10.5441/002/edbt.2021.52.
- [17] S. Shalev-Shwartz, "Online learning and online convex optimization," Foundations and Trends in Machine Learning, vol. 4, no. 2, pp. 107–194, 2012. DOI: 10.1561/22000000018.
- [18] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. Cambridge, MA: MIT Press, 2018.
- [19] J. Zhang, Y. Liu, K. Zhou, G. Li, Z. Xiao, B. Cheng, J. Xing, Y. Wang, T. Cheng, L. Liu, M. Ran, and Z. Li, "An end-to-end automatic cloud database tuning system using deep reinforcement learning," in Proc. ACM SIGMOD Int. Conf. Management of Data, 2019, pp. 415–432. DOI: 10.1145/3299869.3300085.
- [20] S. Chaudhuri and V. R. Narasayya, "AutoAdmin 'what-if' index analysis utility," in Proc. ACM SIGMOD Int. Conf. Management of Data, 1998, pp. 367–378. DOI: 10.1145/276304.276337
- [21] S. Idreos, M. L. Kersten, and S. Manegold, "Database cracking," in Proc. 3rd Biennial Conf. Innovative Data Systems Research, 2007, pp. 68–78.
- [22] A. Deshpande, Z. Ives, and V. Raman, "Adaptive query processing," Foundations and Trends in Databases, vol. 1, no. 1, pp. 1–140, 2007. DOI: 10.1561/19000000001.
- [23] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Humanlevel control through deep reinforcement learning," Nature, vol. 518, no. 7540, pp. 529–533, 2015. DOI:10.1038/nature14236.
- [24] E. S. Page, "Continuous inspection schemes," Biometrika, vol. 41, no. 1/2, pp. 100–115, 1954. DOI: 10.2307/2333009.
- [25] A. Bifet and R. Gavald'a, "Learning from timechanging data with adaptive windowing," in Proc. 2007 SIAM Int. Conf. Data Mining, 2007, pp. 443–448. DOI: 10.1137/1.9781611972771.42.
- [26] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczos, R. Salakhutdinov, and A. Smola, "Deep sets," in Advances in Neural Information Processing Systems, vol. 30, 2017.
- [27] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, "Overcoming catastrophic forgetting in neural networks," Proc. National Academy of Sciences, vol. 114, no. 13, pp. 3521–3526, 2017. DOI: 10.1073/pnas.1611835114.
- [28] Transaction Processing Performance Council, "TPCH Benchmark Specification, Revision 3.0.1," 2023. [Online]. Available: <http://www.tpc.org/tpch/>
- [29] Transaction Processing Performance Council, "TPCDS Benchmark Specification, Version 3.2.0," 2021.