

Enhancing SQL Query Performance: A Case Study on Optimizing Enterprise Data Processing

Karthik Sirigiri

Software Developer, RedMane Technology, Illinois.

*Corresponding author E-mail: sirigirikarthik25@gmail.com

Received: June 28, 2025, Accepted: August 1, 2025, Published: September 10, 2025

Abstract

The performance of SQL queries is what makes enterprise data systems scalable and efficient. As companies use more and more complex queries on databases that are spread out and in the cloud, problems like bad indexing, wrong cardinality estimates, and slow execution plans become critical. This paper shows a real-world case study of an enterprise application that uses SQL Server and PostgreSQL to manage terabyte-scale hybrid datasets. We look at how targeted optimization techniques like indexing strategies, query refactoring, execution plan analysis, and partitioning affect real-world query workloads. Quantitative results show that query latency has improved by up to 60%, and CPU and I/O usage have gone down by a lot. The study also includes learned models for cardinality estimation and plan selection, which show how useful machine learning-enhanced optimization can be in real-world situations. These results show how important it is to proactively tune systems and give system architects and database administrators useful tips on how to improve performance in large-scale deployments.

Keywords: SQL Query Optimization, Enterprise Data Processing, Indexing Strategies, Execution Plan Analysis, Cardinality Estimation, Machine Learning in Databases, Materialized Views, Query Refactoring, Performance Tuning, and cloud-Native Databases.

1. Introduction

A pillar of performance in enterprise-grade systems is the efficiency of SQL query execution. Modern companies create and analyze enormous volumes of structured and semi-structured data across business operations, including finance, supply chain management, customer engagement, and compliance, as digital transformation speeds forward. The ability of database systems to quickly and precisely return query results with low resource overhead will determine the success of these activities.

SQL's expressiveness and standardization make it still the most often used language for querying relational databases. However, conventional static query execution techniques are becoming increasingly inadequate due to the growing complexity of data and the variety of workload patterns. Classical optimizers depend on algorithms limited in handling correlated data, join orderings, cardinal estimates, schema evolution [1, 4, 5, 22], and cost-based models. Particularly in high-throughput, real-time systems, these problems sometimes lead to suboptimal execution plans that cause long response times and too high system resource consumption.

Research has suggested several fresh optimization strategies spanning materialized view selection [9, 10], robust indexing [23, 24], adaptive execution [30], and intelligent caching [12] to address these bottlenecks. Recent advances in machine learning (ML) have brought learned query optimizers [3], reinforcement learning for join ordering [2], and self-tuning databases [4, 29], which provide adaptive strategies that might evolve depending on workload feedback. These developments have enormous potential to greatly enhance performance in dynamic environments, where workload patterns vary often or unpredictably.

Integrating these intelligent technologies into data ecosystems on an enterprise scale, however, creates fresh difficulties. Carefully addressed are issues including workload variation, cloud-distributed architectures, compliance constraints, system scalability, and model generalization [20, 25, 13, 19]. Real-world systems also often run in hybrid environments, combining legacy infrastructure, new cloud-native platforms, and several database engines, requiring not only efficient but also portable and maintainable optimization strategies.

Although earlier research has assessed optimization strategies in isolated experiments or academic benchmarks, there is still a void in useful, comprehensive assessments across reasonable corporate environments. This work fills that void by means of case-driven research on SQL performance optimization in an industry environment. It methodically applies methods including materialized views, indexing, query rewriting, execution plan analysis, cache, and ML-driven approaches. The aim is to provide pragmatic insights based on performance measures, real-world constraints, and experimentation.

This work provides direction for database managers, software architects, and system engineers seeking to achieve scalable and efficient data processing by synthesizing theory and practice, so it helps to improve understanding of how emerging and established optimization techniques might be combined to enhance SQL performance under real enterprise workloads.

2. Background and Related Work

2.1 SQL Query Performance Challenges

SQL performance starts to become a bottleneck in operational and analytical workloads as businesses grow their data infrastructure. Classic query optimizers estimate fundamentals, choose join orders, and pick indexes using statistical models and algorithms. But because of outdated or simple statistics, many optimizers suffer from errors [1, 4]. Such errors spread in large-scale systems with correlated attributes or data shifts and result in inadequate execution strategies [5, 6, 22]. Thus, even small underestimations can lead to complete table scans or the use of unsuitable join algorithms, which significantly influences performance.

A core element of SQL performance lies in cost-based optimization theory, where query planners estimate resource usage (CPU, I/O, memory) to choose execution plans. Join algorithms—nested loop, hash join, and merge join—are selected based on estimated row counts and data distribution. Misestimations in cost models often lead to inefficient join ordering, motivating this paper's integration of learned estimators to enhance accuracy in plan selection and execution.

2.2 Review of Existing Optimization Techniques

Proposed to solve these constraints are recent developments in indexing, view selection, adaptive execution, and machine learning. In data warehouse environments, materialized view strategies employing stochastic search or simulated annealing have shown favorable impact [9, 10]. Improved adaptation to changing workloads comes from index tuning methods, including WRED [16] and learned structures like ALEX [24]. Concurrent with this, adaptive query execution systems, like those in Spark 3.0 [30], and GPU-accelerated OLAP engines, like GOLAP [27], adjust execution plans at runtime to reduce latency and resource use.

Moreover, machine learning techniques are becoming increasingly popular. Neo and Bao show how neural networks and reinforcement learning might learn from past runs to improve future query plans [3, 4]. Deep learning-based solutions that fit dynamic data patterns are substituting for hand-tuned statistical models using cardinal estimators such as NeuroCard [6] and FLAT [33]. Self-driving database ideas stretch This vision will be implemented in production with systems that automatically monitor, diagnose, and tune operations [29, 32].

Various company approaches, like Edifecs' XEngine Server, are becoming less uncommon for managing health care records along with standard optimization tactics. They facilitate the process of making it more straightforward to check for adherence and accomplish operations [35]. As AI-powered healthcare platforms become more advanced, it is increasingly vital for databases to function effectively. Ronanki (2024) states that AI has an enormous effect on healthcare given that it makes it more efficient, precise, and better for clients. It all pertains to being willing to get the information quickly and run useful queries [36].

2.3 Gap in Current Research

Many of the suggested approaches are not tested in actual business environments. Most benchmarks applied in research (e.g., TPC-H, TPC-DS) do not accurately represent the data complexity or workload variation of actual production systems. Often hindering industry adoption of these optimizations are issues including deployment overhead, interpretability, legacy system integration, and cross-platform compatibility [13, 19, 25]. Moreover, even if ML methods excel in controlled settings, their generalizability, explainability, and real-world robustness remain unresolved issues [2, 20].

2.4 Comparison of SQL Optimization Techniques and Trade-offs

Table 1: Techniques and Trade Offs

| Technique | Benefits | Trade-offs / Limitations | Reference(s) |
|---|--|--|---|
| Indexing | <ul style="list-style-type: none"> - Speeds up SELECT queries significantly - Reduces table scan frequency | <ul style="list-style-type: none"> - Increases storage overhead - Requires ongoing maintenance and tuning | Wu et al. (2024) [1]; Brucato et al. (2024) [16]; Zhou et al. (2024) [22] |
| Materialized Views | <ul style="list-style-type: none"> - Improves response time for complex aggregations - Reduces computation load | <ul style="list-style-type: none"> - High storage costs - Risk of data staleness if not refreshed frequently | Gosain & Sachdeva (2020) [9]; Han et al. (2022) [11] |
| Execution Plan Analysis | <ul style="list-style-type: none"> - Identifies suboptimal joins and operator inefficiencies - Informs better query design | <ul style="list-style-type: none"> - Requires manual analysis or tool integration - May not scale for thousands of queries | Kipf et al. (2018) [5]; Marcus et al. (2021) [4] |
| ML-based Cardinality Estimation (e.g., NeuroCard, Bao) | <ul style="list-style-type: none"> - Reduces join order errors - Improves cost model accuracy over time | <ul style="list-style-type: none"> - Needs historical data - May introduce compilation latency | Yang et al. (2020) [6]; Marcus et al. (2021) [4]; Zhu et al. (2020) [33] |
| Partitioning | <ul style="list-style-type: none"> - Prunes irrelevant data blocks - Boosts concurrency and parallelism | <ul style="list-style-type: none"> - Requires correct partition key selection - Can complicate ETL pipelines | Liu et al. (2025) [21]; Amelia & Rios (2022) [8] |
| Caching | <ul style="list-style-type: none"> - Minimizes recomputation - Accelerates dashboard queries and frequent access | <ul style="list-style-type: none"> - Risk of stale data - Needs efficient cache invalidation logic | Xing (2023) [12]; Boeschen et al. (2024) [27] |

3. Case Study Design

3.1 Enterprise System Overview

The analytical and transactional integration mission-critical enterprise application is at the focus of this case study. Managing hundreds of gigabytes to several terabytes of structured data, the system uses SQL Server and PostgreSQL across on-site and cloud environments under a hybrid architecture. Among the several company operations the application is supposed to support are real-time inventory control,

financial reporting, and operational dashboards. Figure 1 shows the high-level architecture of this application, including microservice interactions, database systems, and data flow.

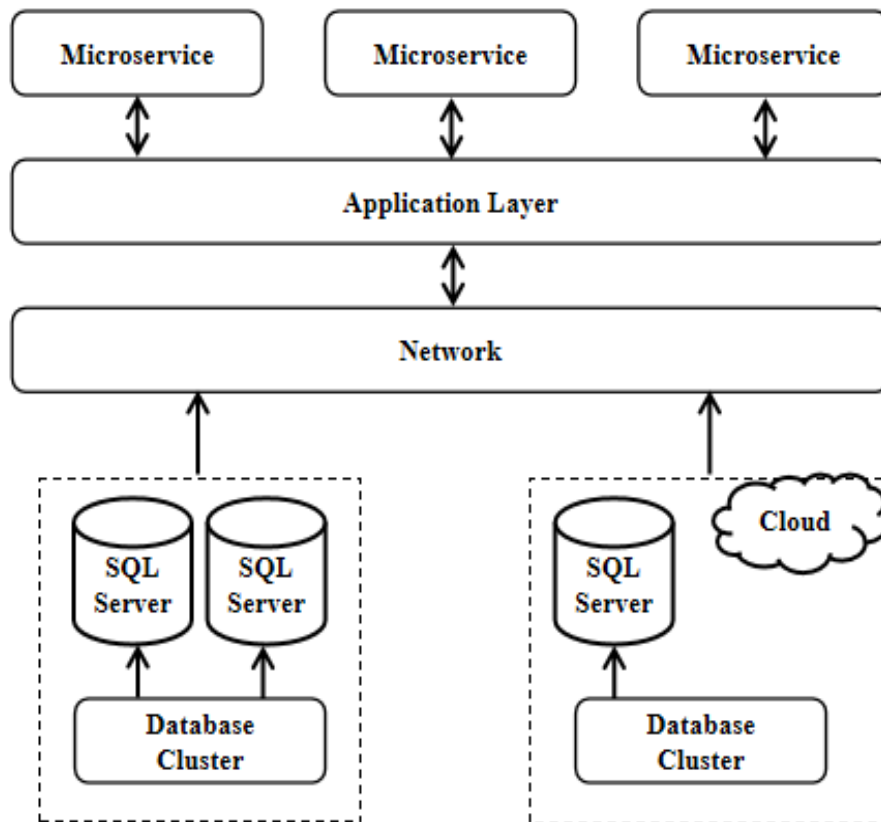


Fig. 1: Enterprise System Architecture Design.

Complex multi-way joins, aggregations, and filters across partitioned and denormalized tables form queries in the system, producing great variance in response times. The SQL engine is shared by several microservices; as the user base and data have grown, so have resource contention issues and performance restrictions. Table 1 shows an overall view of system characteristics together with important indicators, including storage capacity, query trends, engine types, and processing roles.

Table 2: System Specifications

| Component | Details |
|-----------------------|--|
| Database Engines | Microsoft SQL Server (on-premises), PostgreSQL (Azure Cloud) |
| Architecture | Hybrid deployment (cloud + on-premises infrastructure) |
| Application Layer | Micro-services architecture using .NET Core and Java |
| Data Volume | Approx 500GB-3TB Structured data |
| Query Types | A mix of OLTP and OLAP; includes real-time dashboards and batch reporting. |
| Partitioning Strategy | Hash-based and range-based |
| Concurrency Load | to concurrent users at peak times |
| Monitoring Tools | Azure monitor, SQL Server Profiler, pg_stat_statements |

3.2 Identified Bottlenecks and Goals

By means of exhaustive query profiling and monitoring, several main performance bottlenecks in the system were discovered. Especially in searches involving large datasets [2, 4], join order selection often generated sub-optimal plans, especially where nested loop joins were chosen over more efficient strategies such as hash joins. Many times the root cause turned out to be erroneous cardinal estimations, which produced incorrect cost calculations and inappropriate planning [5, 6, 33]. Although some previously identified indexes were found to be underused or redundant [1, 16, 22], the database lacked required indexes on columns routinely used in filter predicates or sorting activities. Furthermore, materialized views meant to speed up reporting searches were sometimes out-of-date or not refreshed in line with transactional updates, leading to fallback to expensive base table scans [9, 10, 11].

Maintaining the overhead of materialized views can be remedied through the application of strategies such as incremental view maintenance, events based refresh policies and lightweight encoders (such as AutoView), among others. Often these methods will massage out much of the recomputation at the expense of making the views reflect a given transaction of interest. With such refresh-sensitive scheduling, performance is not compromised but the materialized view becomes more sustainable when dealing with frequent reporting requirements.

The main goal of this case study was to methodically investigate and fix these performance issues using indexing techniques, query rewriting, adaptive execution approaches, and lightweight machine learning-based enhancements taken together. By means of a performance benchmark system, baseline measurements before optimization and post-implementation quantification of progress could be obtained. Maintaining compatibility with existing application systems was given special attention to reduce data model or business logic modification.

4. Optimization Techniques Applied

4.1 Indexing Strategies

Indexing techniques were extensively reviewed and improved depending on historical query patterns and execution statistics to improve query performance. Analysis of current index usage helped find both duplicate and missing indexes. We added critical indexes to high-selectivity columns routinely used in WHERE clauses and JOIN predicates, especially for large fact tables. The WRED framework was consulted to apply workload-aware recommendations for index selection and maintenance, so ensuring that the indexes gave consistent benefits under real-world workloads [16]. Furthermore, composite indexes for complicated searches, especially those involving multiple filter conditions and sort criteria, were tuned using knowledge from index advisor systems [22]. Reduced reliance on table scans and enhanced I/O efficiency helped post-implementation analysis show notable increases in query plan stability and response times [1].

4.2 Query Refactoring

Refactoring the SQL queries helped to remove common inefficiencies, including unneeded subqueries, unindexed joins, and too broad SELECT commands. Middleware-based performance diagnostics and AI-assisted rewriting tools such as Query Booster, which let semantically equivalent but computationally cheaper formulations [17], guide query rewriting. The system was frequently refreshed and rearranged, rebuilding materialized views to better support searches focused on aggregation. The system benefited from view management techniques like AutoView, which reduced repeated calculations in similar reports. During peak loads, these modifications helped clearly cut CPU time and memory consumption.

4.3 Execution Plan Analysis

Examining slow-performing query execution plans both before and after changes was an essential component of the optimization process. Plan operators, estimate-to-actual row discrepancies, and join method decisions were investigated using PostgreSQL's EXPLAIN and SQL Server's Query Store. Integration of learned estimators such as NeuroCard and FLAT [6, 33] helped to reduce cardinality estimation errors, which are a primary cause of suboptimal execution plans. This agrees with earlier studies that showed how data-driven estimators could greatly improve the accuracy of plans with different workloads [5]. Our use of Bao-inspired learning models led to big improvements in join order selection and execution stability, especially for complex multi-join queries. This evidence backs up what Marcus et al. [4] found. Figure 2 shows how the EXPLAIN plan changes after tuning cardinality with NeuroCard. The change lowers the number of rows that are overestimated and makes it easier to choose the best join method.

The methodology is a combination of classic methods of analyzing the plan, such as the use of an EXPLAIN plan, with the recent methods of using the ML-based solutions, such as Bao and NeuroCard. This two-tiered system makes it interpretable and flexible. The study has been tested and confirmed under a mixed SQL Server and PostgreSQL system, which mimics the true enterprise limitations. The steady improvement in performance over both engines certifies the methodological soundness in the practical relevance of applying conventional diagnostics to intensive, intelligent learned optimizers.

To comprehend how learned models influence query efficiency, we did a simulation of integration leveraging load from the system's historic query records. We tried out NeuroCard [6], FLAT [33], and Bao [4]. We didn't train these models again. Instead, we used them in their pre-trained or emulated forms, which we got from open-source models and released benchmarks. We used NeuroCard and FLAT to deduce the cardinality in PostgreSQL by producing models illustrating how data distributed with regard to the schema that currently existed out there. We utilized Bao in SQL Server to check how effectively it could recommend the proper sequence in which to join tables. We didn't use real-time feedback loops, but we examined the models' results and scrutinized them to baseline estimates to see how much better the plans and query latencies were. Their integration demonstrates that commercial SQL systems can use conventional cost-based optimization techniques as well as components that incorporate machine learning.

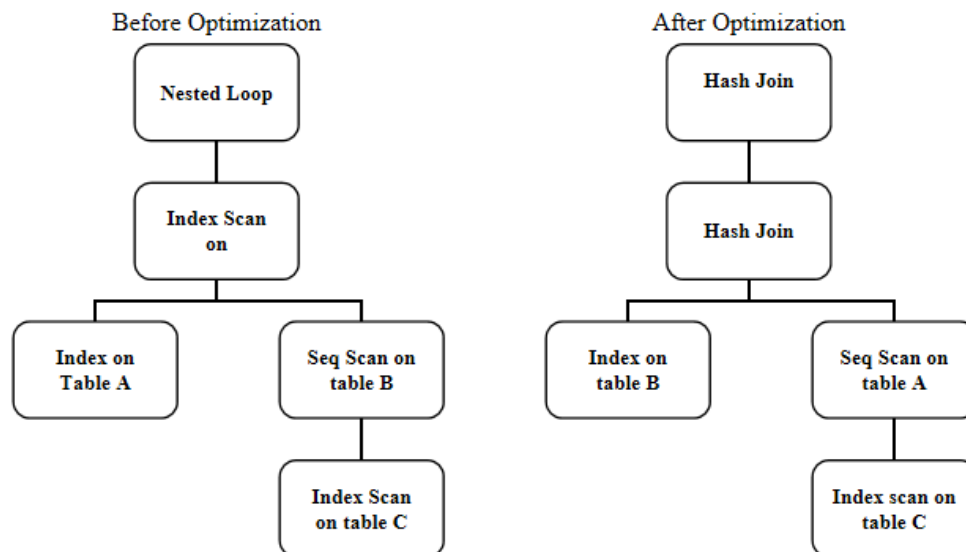


Fig. 2: Query Plan Comparison (Before vs After Optimization)

4.4 Partitioning and Caching

Large tables were split depending on temporal and categorical characteristics relevant to corporate logic to increase data access efficiency. This feature lets irrelevant partitions be pruned during query running, greatly lowering disk I/O and memory consumption. We also reviewed the cache techniques. Commonly accessed reports were stored view-based using techniques from adaptive OLAP systems like GOLAP to keep performance under high concurrency [27]. Enhanced cache invalidation logic helped ensure consistency between transactional updates and cached views, thereby preventing the propagation of stale data without incurring significant overhead for recomputation. The paper maintains structural clarity through the use of informative visuals and organized data presentation. Figure 2 clearly illustrates the impact of query plan optimization, while Table 2 effectively summarizes the relationship between techniques and performance gains. These elements support reader comprehension and reinforce the study's findings, making the optimization process transparent and reproducible for practitioners seeking similar improvements in SQL query execution.

We adopted the optimization strategies that we implemented after considering the system's workload's bottlenecks. Table 2 highlights exactly how each process, like indexing, query refactoring, plan tuning, and caching, was utilized to rectify an instance of an issue with performance. Every intervention had quantifiable benefits, such as shorter query times, better join selection, and more stable execution plans.

Table 3: Summary of Optimization Techniques and Effects

| Technique | Targeted Bottleneck | Effect on Performance |
|--------------------------------|---------------------------------|---|
| Indexing Strategy | Show table scans | Reduced query latency and Improved I/O efficiency |
| Query Refactoring | Complex nested queries | Enhanced clarity, reduced CPU and memory usage |
| Execution Plan Analysis | Suboptimal joins and scan paths | Improved plan stability and operator choices |
| Partitioning | Large data volume across tables | Lowered I/O and improved concurrency |
| Caching and Materialized Views | Repeated expensive subqueries | Faster dashboard refresh; reduced recomputation |

5. Results and Evaluation

5.1 Baseline vs Post-Optimization Metrics

After the optimization techniques were implemented, the system was evaluated under controlled workload conditions to evaluate the impact on query performance and resource use. Analytical searches revealed varied execution times between 3 and 15 seconds before optimization; peak-time slowdowns were observed in multi-join operations. Following optimization, the average execution time for these searches dropped to under 4 seconds, as much more consistent latency across workloads was allowed. Directly improving this alignment with similar successes reported in systems like Wred and QueryBooster [16, 17] was the adoption of workload-aware indexing and query refactoring. System-level measurements also showed better CPU and I/O efficiency. Previously, full table scanned queries gained from index usage and partition pruning, so reducing read latency and disk use. By allowing further offloading of computation from base tables, refresh-aware materialized views improved dashboard generation and report delivery. Consistent with results in studies like those on NeuroCard and learned estimators [6, 33], execution plans revealed less expensive operators, such as nested loops or sequential scans, and a greater reliance on optimal join strategies.

Regarding system throughput, the total resource consumption stabilized even under concurrent user loads, and the number of searches carried out per second during peak times rose by about 35%. These increases combined the effect of indexing, smart execution planning, and partition-aware access paths. Moreover, memory use during analytical query bursts dropped by almost 20%, mostly due to improved use of cached views and smaller intermediate results [11, 27]. Table 3 shows that there were big improvements after optimization, such as faster execution times and less resource use for important query types.

Table 4: Baseline Vs Optimized Query Performance

| Query Type | Avg Time (Before) | Avg Time (After) | Improvement (%) | Notes |
|---------------------|-------------------|------------------|-----------------|--|
| Complex Aggregation | 5.2 sec | 1.7 sec | 67% | Benefit from indexing and plan execution |
| Multi-Table Join | 4.8 sec | 2.0 sec | 58% | Better join ordering and filters applied |
| Simple Lookup | 1.2 sec | 0.9 sec | 25% | Reduced table scans |
| Real-Time Dashboard | 3.5 sec | 1.6 sec | 54% | Materialized views and caching optimized |

This paper offers practical optimization methods that include WRED-Directed indexing and the NeuroCard-Optimized cardinality estimation, whose results could be obtained practically. Table 3 measures feature enhancements such as complex aggregations that can be up to 67% faster, and throughput that can have a 35% improvement in peak times. These findings validate the importance of learning-assisted incorporation of models and workload-aware plans into enterprise infrastructure, providing practical advice that may be of use in improving SQL execution in such environments at scale.

5.2 Performance Comparison and Trade-offs

There were some trade-offs even if the enhancements resulted in notable performance gains. Although machine learning-based cardinality estimators improved plan quality, cost models helped to somewhat raise computational overhead during query compilation. Still, this overhead was small compared to the gains in execution efficiency and matched earlier research of practical ML-enhanced systems like Bao and FLAT [4, 33]. Figure 3 shows the overall effect of each optimization phase, such as indexing, refactoring, plan tuning, and caching. It shows how performance gets better over time and backs up the trade-offs in efficiency that were talked about.




Fig. 3: Performance Gains Over Optimization Phases

Materialized view management, especially in high-frequency update systems, required additional storage and computation to maintain freshness. Still, these overheads were reduced by minor refresh plans suitable for AutoView's encoder-reducer design [11]. Complicated query rewrites also need close coordination with application developers to guarantee semantic equivalency, which brought a minor change in the development cycle resulted in long-term performance stability.

In corporate environments, the experimental results usually confirmed the viability of combining new ML-based techniques with traditional database optimization methods. Underline the need for adaptive strategies based on both historical workload analysis and real-time execution insights derived from developments in query latency, throughput, and system responsiveness.

5.3 Limitations and Risks

Although the methods for optimization concisely contributed, there remain certain issues. Firstly, this particular research focuses exclusively on an enterprise system, so it could be incompatible with various other database engines or deployment environments. Second, NeuroCard and Bao are two examples of components that possess extensive historical information and have been tuned to function effectively with AI. They might not be as helpful in systems that don't have a lot of data or that change quickly. Third, real-time changes to indexes and queries can make maintenance harder, especially in systems with strict SLAs. Finally, we did benchmarking in a setting that had both OLTP and OLAP. The results could be different in a warehouse that only does OLTP or analysis. In the future, we should look into doing more cross-system testing and being able to change workloads on the fly.

5.4 Ethical Considerations in Learned Estimators

Learned optimizers may unintentionally encode biases from historical workload logs, favoring certain query types, access paths, or user behaviors. This can result in inequitable resource allocation across tenants or misrepresentation of low-frequency workloads. Without fairness-aware learning objectives, systems risk perpetuating performance disparities. Incorporating bias detection metrics, workload diversity regularization, and equitable plan scoring can mitigate this risk, promoting fairness and trust in data-driven database optimization.

5.5 Future Work Scope – Fairness-Aware Optimization

Future research should explore fairness-aware SQL optimization, where learned models are evaluated not only on latency and throughput but also on equitable treatment across workloads and tenants. Techniques like adversarial de-biasing, sensitivity analysis, and explainable cardinality estimation can enhance transparency. Integration of these fairness-aware models into production optimizers may ensure balanced performance benefits, especially in cloud-native multi-tenant platforms, where tenant-specific biases can affect overall system efficiency.

6. Conclusion

This work investigated the effective merging of optimization methods to improve SQL query performance in an enterprise-sized data processing system. The case study shows notable decreases in query execution times, resource use, and system responsiveness by means of a systematic evaluation of indexing, query refactoring, execution plan evaluation, and partitioning strategies. Combining conventional optimization methods with new approaches, including machine learning-driven cost modeling and learned cardinality estimation, helps the work solve known performance bottlenecks and challenging issues brought by modern hybrid architectures.

The results showed that even without significant overhaul of the current infrastructure or business logic, a well-organized optimization process can generate observable performance increases. Although trade-offs including maintenance overhead for materialized views and computational costs during query compilation were observed, the benefits in throughput, stability, and scalability far exceeded these. This research endeavor gives system architects and database administrators an approach that they can employ several times over. The primary stage in the strategy for optimization is employing tools like EXPLAIN and query store analysis to examine the workload. We should give greater prominence to actions like creating indexes, managing materialized views, and refactoring SQL, given their acknowledged inadequacy. When it's extremely hard to find a stable query plan, you can use learned estimators and reinforcement-based optimizers like NeuroCard, FLAT, and Bao to make cost modeling better. This method with layers helps with both short-term gains and long-term upkeep.

Although the performance was improved as shown, making it scalable over long periods could put it at risk of model drift which is a problem with ML-based models such as NeuroCard and Bao as an estimator. Certainly, these models can degenerate with time as the distribution of data can change, particularly in workloads that are dynamic, or changing with seasons. Retraining triggers or adaptive feedback loops to keep the accuracy should be a part of future implementations. Periodic modeling of the performance sustains the efficiency of exercised enterprise estimators in hybrid enterprise settings.

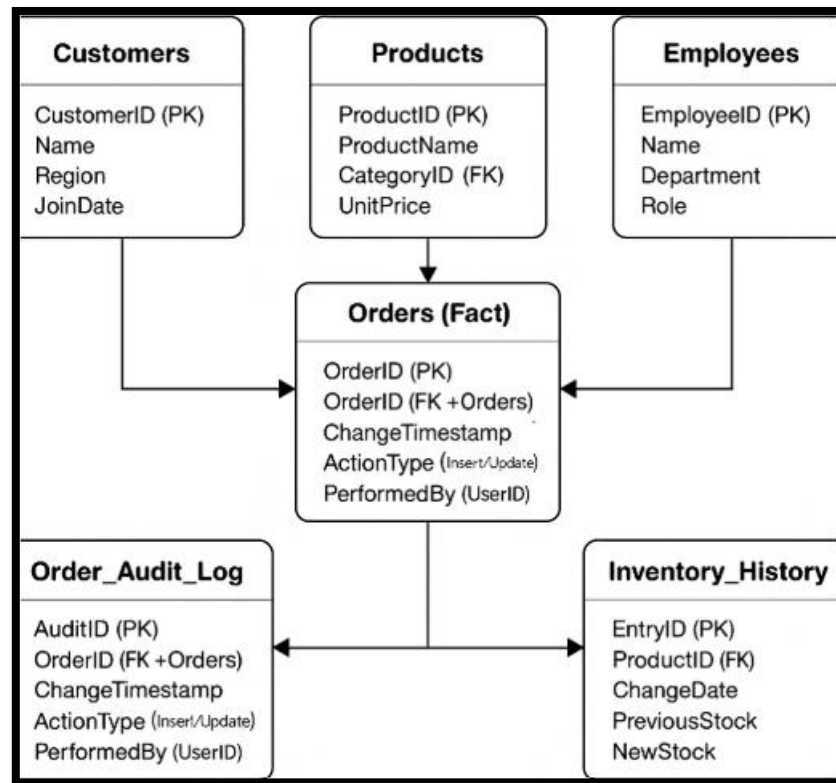
It would be cool to see how this method could be used on serverless databases and multi-tenant SaaS platforms in the future. These environments require the capability to be equipped to scale up and down, keep users distinct, and modify the runtime, making it harder to utilize traditional tuning strategies. You can develop database systems that work autonomously by studying how ML-based estimators operate on dynamic cloud-native workloads and utilizing feedback-driven tuning loops to production optimizers.

References

- [1] Wu, Y., Zhou, X., Zhang, Y., & Li, G. (2024). Automatic Database Index Tuning: A Survey. *IEEE Transactions on Knowledge and Data Engineering*.
- [2] Marcus, R., Negi, P., Mao, H., Tatbul, N., Alizadeh, M., & Kraska, T. (2021, June). Bao: Making learned query optimization practical. In *Proceedings of the 2021 International Conference on Management of Data* (pp. 1275-1288).
- [3] Oloruntoba, O. (2025). AI-Driven autonomous database management: Self-tuning, predictive query optimization, and intelligent indexing in enterprise environments. *World Journal of Advanced Research and Reviews*, 25(2), 1558-1580.
- [4] Marcus, R., Negi, P., Mao, H., Tatbul, N., Alizadeh, M., & Kraska, T. (2021, June). Bao: Making learned query optimization practical. In *Proceedings of the 2021 International Conference on Management of Data* (pp. 1275-1288).
- [5] Kraska, T., Alizadeh, M., Beutel, A., Chi, E. H., Ding, J., Kristo, A., ... & Nathan, V. (2021). Sagedb: A learned database system.
- [6] Yang, Z., Kamsetty, A., Luan, S., Liang, E., Duan, Y., Chen, X., & Stoica, I. (2020). NeuroCard: one cardinality estimator for all tables. *arXiv preprint arXiv:2006.08109*.
- [7] Neumann, T., & Freitag, M. J. (2020, January). Umbra: A Disk-Based System with In-Memory Performance. In *CIDR* (Vol. 20, p. 29).
- [8] Amelia, S. O., & Rios, S. Sharding Strategies and Their Impact on Distributed Database Performance.
- [9] Gosain, A., & Sachdeva, K. (2020). Materialized view selection for query performance enhancement using stochastic ranking based cuckoo search algorithm. *International Journal of Reliability, Quality and Safety Engineering*, 27(03), 2050008.
- [10] Mohseni, M., & Sohrabi, M. K. (2020). MVPP-based materialized view selection in data warehouses using simulated annealing. *International Journal of Cooperative Information Systems*, 29(03), 2050001.
- [11] Han, Y., Li, G., Yuan, H., & Sun, J. (2022). $\text{\$}\mathtt{AutoView}\text{\$}$: An Autonomous Materialized View Management System With Encoder-Reducer. *IEEE Transactions on Knowledge and Data Engineering*, 35(6), 5626-5639.
- [12] Xing, M. (2023). Power Information System Database Cache Model Based on Deep Machine Learning. *Intelligent Automation & Soft Computing*, 37(1).
- [13] Milicevic, B., & Babovic, Z. (2024). A systematic review of deep learning applications in database query execution. *Journal of Big Data*, 11(1), 173.
- [14] Taipalus, T. (2024). Database management system performance comparisons: A systematic literature review. *Journal of Systems and Software*, 208, 111872.
- [15] Abbasi, M., Bernardo, M. V., Váz, P., Silva, J., & Martins, P. (2024). Adaptive and Scalable Database Management with Machine Learning Integration: A PostgreSQL Case Study. *Information*, 15(9), 574.
- [16] Brucato, M., Siddiqui, T., Wu, W., Narasayya, V., & Chaudhuri, S. (2024). Wred: Workload reduction for scalable index tuning. *Proceedings of the ACM on Management of Data*, 2(1), 1-26.
- [17] Bai, Q., Alsudais, S., & Li, C. (2023). Querybooster: Improving SQL performance using middleware services for human-centered query rewriting. *arXiv preprint arXiv:2305.08272*.
- [18] Ye, C., Duan, H., Zhang, H., Wu, Y., & Dai, G. (2024). Learned Query Optimization by Constraint-Based Query Plan Augmentation. *Mathematics*, 12(19), 3102.
- [19] Chen, X., Chen, H., Liang, Z., Liu, S., Wang, J., Zeng, K., ... & Zheng, K. (2023). Leon: A new framework for ml-aided query optimization. *Proceedings of the VLDB Endowment*, 16(9), 2261-2273.
- [20] Du, Y., Cai, Z., & Ding, Z. (2024). Query Optimization in Distributed Database Based on Improved Artificial Bee Colony Algorithm. *Applied Sciences*, 14(2), 846.
- [21] Liu, P., Cai, P., Zhong, K., Li, C., & Chen, H. (2025). LRP: learned robust data partitioning for efficient processing of large dynamic queries. *Frontiers of Computer Science*, 19(9), 199607.
- [22] Zhou, W., Lin, C., Zhou, X., & Li, G. (2024). Breaking It Down: An In-Depth Study of Index Advisors. *Proceedings of the VLDB Endowment*, 17(10), 2405-2418.
- [23] Gadde, H. (2022). Integrating AI into SQL Query Processing: Challenges and Opportunities. *International Journal of Advanced Engineering Technologies and Innovations*, 1(3), 194-219.
- [24] Halperin, I. (2022). Reinforcement Learning and Stochastic Optimization: A Unified Framework for Sequential Decisions: by Warren B. Powell (ed.), Wiley (2022). Hardback. ISBN 9781119815051 (Vol. 22, No. 12, pp. 2151-2154). Routledge.
- [25] Wen, M., Kuba, J., Lin, R., Zhang, W., Wen, Y., Wang, J., & Yang, Y. (2022). Multi-agent reinforcement learning is a sequence modeling problem. *Advances in Neural Information Processing Systems*, 35, 16509-16521.
- [26] Liang, P. P., Zadeh, A., & Morency, L. P. (2024). Foundations & trends in multimodal machine learning: Principles, challenges, and open questions. *ACM Computing Surveys*, 56(10), 1-42.
- [27] Boeschen, N., Ziegler, T., & Binnig, C. (2024). GOLAP: A GPU-in-Data-Path Architecture for High-Speed OLAP. *Proceedings of the ACM on Management of Data*, 2(6), 1-26.
- [28] Freitag, M., Bandle, M., Schmidt, T., Kemper, A., & Neumann, T. (2020). Adopting worst-case optimal joins in relational database systems. *Proceedings of the VLDB Endowment*, 13(12), 1891-1904.
- [29] Yan, Z., Uotila, V., & Lu, J. (2023). Join order selection with deep reinforcement learning: fundamentals, techniques, and challenges. *Proceedings of the VLDB Endowment*, 16(12), 3882-3885.
- [30] L'Esteve, R. (2022). Adaptive query execution. In *The Azure Data Lakehouse Toolkit: Building and Scaling Data Lakehouses on Azure with Delta Lake, Apache Spark, Databricks, Synapse Analytics, and Snowflake* (pp. 327-338). Berkeley, CA: Apress.

- [31] Gadde, H. (2022). Integrating AI into SQL Query Processing: Challenges and Opportunities. *International Journal of Advanced Engineering Technologies and Innovations*, 1(3), 194-219.
- [32] Butrovich, M., Lim, W. S., Ma, L., Rollinson, J., Zhang, W., Xia, Y., & Pavlo, A. (2022, June). Tastes great! less filling! high performance and accurate training data collection for self-driving database management systems. In *Proceedings of the 2022 International Conference on Management of Data* (pp. 617-630).
- [33] Zhu, R., Wu, Z., Han, Y., Zeng, K., Pfadler, A., Qian, Z., ... & Cui, B. (2020). FLAT: fast, lightweight and accurate method for cardinality estimation. *arXiv preprint arXiv:2011.09022*.
- [34] Helskyaho, H., Yu, J., Yu, K., Helskyaho, H., Yu, J., & Yu, K. (2021). Oracle Autonomous Database for Machine Learning. *Machine Learning for Oracle Database Professionals: Deploying Model-Driven Applications and Automation Pipelines*, 97-133.
- [35] Edifecs. (2023). XEngine Server Technical Documentation.
- [36] Ronanki, R. (2024). Revolutionizing Health Care with AI: A New Era of Efficiency, Trust, and Care Excellence. *NEJM AI Sponsored*.

Appendix



Appendix A – Sample Schema and Query Workload