

Review Article

Review of Query Optimization Techniques in Database Management Systems

Kapil Ahir^{1*}

Assistant Professor, Mandsaur University, Mandsaur, Department of Computer Sciences and Applications, India

Received 01 May 2026, Accepted 18 May 2026, Available online 19 May 2026, Vol.16, No.3 (May/Jun 2026)

Abstract

In database management systems (DBMSs), query optimization is a frequent practice. Its primary goal is to maximize query execution efficiency among many possibilities. In recent years, the world of data-driven applications has grown rapidly, and effective query processing is an essential element of high-performance systems and efficient resource use in cloud computing environments. In this paper, the authors review key query optimization strategies used in state-of-the-art DBMSs. These include heuristic-based optimization, cost-based optimization, semantic-based optimization for both join and indexing strategies and sub-query rewriting methods. The study contextually explores new AI/ML-based optimization techniques that contribute towards enhancing adaptive query execution and automated tuning in the dynamic environment. Apart from this, the paper also discusses the role of query execution plans, cost models and optimization frameworks and how they can improve the performance of databases in enterprise systems, databases in a distributed environment of the cloud and big data analytics platforms. Using a systematic literature review of recent relevant studies, researchers identify current progress and challenges in query optimization alongside possible future research directions. The paper shows that in the context of many complex (possibly correlated) workloads, modern database systems need some degree of intelligent/elastic behaviour to operate with bounded latency and resource consumption.

Keywords: Database Management Systems (DBMS), Query Optimization, Query Processing, Distributed Databases, Big Data Analytics, Cloud Databases.

Introduction

Large amounts of data can be efficiently stored, retrieved, and managed thanks to database management systems (DBMSs), which are essential in today's computer environment. The swift expansion of digital apps, cloud computing, e-commerce platforms, and social networking systems [1][2], and big data technologies, databases are required to process increasingly complex queries within a minimal response time [3][4]. In such environments, query optimization has become one of the most crucial elements of database systems as ineffective query execution can result in high memory usage, increased CPU utilization, high disk input/output operations, and poor overall system performance. The process of choosing the best execution method for a particular query from a range of options is known as query optimization. A query can typically be optimized to run using several different logically equivalent execution plans [5].

The goal of the query optimizer is to select one of these execution plans that has the lowest estimated cost of execution, yet still remains correct and reliable.

Previously, conventional query optimization methods were mainly based on heuristic and cost-based optimization methods. Heuristic optimization uses a set of rules to combine its queries and smoothen the intermediate results, and cost-based optimization assumes a large number of query execution plans using layout statistical information and formative cost designs [6]. To boost the performance of a large-scale database system, modern database system has increasingly adopted a number of advanced techniques, including semantic optimization, adaptive query processing, indexing, join optimization, and query rewriting techniques. In addition, new technologies such as Artificial Intelligence (AI) and Machine Learning (ML) are transforming today's query optimization by introducing intelligent cardinality estimation, adaptive execution planning, workload prediction and automated database tuning.

Query optimization is a major aspect that applies to many application domains, such as enterprise database systems, distributed databases and cloud computing

*Corresponding author's ORCID ID: 0000-0000-0000-0000
DOI: <https://doi.org/10.14741/ijcet/v.16.3.1>

platforms [7][8], real-time analytics systems and data warehouses. Optimization methods can help to enhance the execution time and scalability, decrease the resource consumption, and process high-performance transactions [9]. Modern database systems face some challenges, however, with diverse workloads, different data sources (structured vs unstructured) and distributed architecture, such as inaccurate cardinality estimation, query optimization remains a very complex problem. Thus, ongoing research and innovations must create modern intelligent and adaptive optimization frameworks for present-day data processing requirements. This paper surveys the query optimization techniques on DBMSs, including their principles, benefits, drawbacks and applications, along with a review of some emerging research directions.

Structure of the Paper

This paper is organized into several sections. Section II presents an overview of DBMS query processing. Section III discusses various query optimization techniques. Section IV explains applications, challenges, and emerging trends in query optimization. Section V provides a detailed literature review of recent studies. Finally, Section VI concludes the paper and highlights future research directions.

Overview of DBMS Query Processing

The query processing is the process of transforming a high-level SQL query into an efficient low-level execution plan. The goal of query processing is to accurately carry out database operations at the lowest possible cost while improving system performance [10].

Phases of Query Processing

In a typical DBMS, there are four major phases in total query processing, namely, parsing, query rewrite, optimization and execution.

Parsing and Translation

The first one breaks up the SQL statement to verify its structure and meaning. The parser checks the validity of the query based on SQL rules and whether the tables, attributes, and operations used in the query are actually in the database schema. Once validated, the query is transformed to an internal representation, which may be a parse tree, a relational algebra expression, etc.

For example:

```
SELECT Name FROM Employee WHERE Salary > 50000;
```

The parser extracts from the query relation, attributes and selection conditions.

Query Rewriting

The second stage of query rewriting is performing transformations to create logically equivalent but more performance-friendly statements. In this process, redundant conditions are eliminated, and relational algebra rules such as flattening nested subqueries can be applied. Query rewriting helps reduce intermediate results and improves execution efficiency.

Query Optimization

Based on specific metrics, the query optimizer generates multiple execution plans and selects one with the lowest estimated cost. Optimizer is responsible for determining access methods, join orders, indexing strategies and execution operators.

Query Execution

Finally, the execution engine executes the selected query execution plan using physical operators such as scans, joins, sorting operations and aggregation functions. Query execution plays a crucial role in processing massive-scale data and also in high-performance database operations.

Query Execution Plans

The Query Execution Plan (QEP) comprises the logical sequence of physical operations chosen by the optimizer to perform a database query. Execution plans dictate how data is read, processed and joined to produce the output.

The same query may have many execution strategies. A query plan's effectiveness is influenced by a number of variables, such as data distribution, join algorithms, indexing methods, and memory availability.

Common operations included in query execution plans are:

- Sequential table scans
- Index scans
- Nested loop joins
- Hash joins
- Sort-merge joins
- Aggregation operations
- Sorting and grouping operations

For example, a query involving multiple tables may be executed using different join orders. In general, joining smaller relations first increases performance and decreases intermediate result sizes.

DBMSs in modern times use advanced optimization methods to create execution plans without requiring human input. Database administrators can use

execution plans to find performance issues which they can then use to enhance query execution speed.

Cost Models and Optimization

Cost models serve as mathematical tools that query optimizers use to determine resource needs for various execution plans [11]. The plan with the lowest anticipated execution cost is chosen by the optimizer after comparing predicted expenses.

The major optimization objectives include:

- Minimizing query execution time
- Reducing CPU utilization
- Minimizing memory consumption
- Reducing disk input/output operations
- Minimizing network communication overhead
- Improving scalability and throughput

Cost estimation is based on statistical information maintained by the database system. Common statistics include:

- Table cardinality
- Number of tuples
- Data distribution
- Histogram statistics
- Index selectivity
- Page access cost

The current cost models become more intricate because they need to assess distributed systems together with cloud infrastructure, parallel processing systems and diverse workload types [12]. The process of cost estimation needs to be precise because faulty estimates create execution plans that operate ineffectively, leading to diminished system performance.

Query Optimization Techniques in Database Management Systems

The DBMS carries out a heavy-lifting function called query optimization. It is the process of selecting among a number of comparable options the best strategy for carrying out a query [13]. A single SQL query can often be executed by many different but logically equivalent execution plans, so the challenge for an optimizer is to select a plan with low execution cost while also executing it correctly and efficiently. Query optimization leads to faster response times, reduced resource usage and improved scalability of the database system.

Heuristic-Based Optimization

In query optimization, heuristic-based optimizations, also known as rule-based optimizations, were among the earliest. The method relies on hardcoded rules or heuristics to map a query to an improved execution

plan, rather than performing an extensive cost evaluation. Its goal: logical transforms that help improve the cost of executing a query.

A widely-known heuristic for this is selection pushdown, meaning a basic idea by which the database can apply selections as early as possible. This not just reduces the size of intermediate results, it also filters unnecessary tuples at an early stage. In a similar fashion, projection pushdown minimizes the number of attributes processed by only reading necessary columns at the start of execution.

Another important heuristic technique is join reordering because it operates by performing smaller relation joins before executing larger relation joins, which results in decreased size of intermediate results. The system uses three additional heuristics: eliminating duplicate processes, reducing complex logical expressions and converting nested queries into their equivalent join forms.

Heuristic optimization provides multiple benefits through its straightforward design and minimal optimization requirements and quick execution plan development. The method achieves computational efficiency because the optimizer limits its assessment of execution plans to a single plan. The method proves effective for both small databases and medium-sized databases [14].

Heuristic optimization methods function effectively until they reach their boundary limitations. The selected execution plan will not achieve global optimality because the method fails to assess both real data statistics and execution costs. Heuristic methods generate inefficient execution plans when handling complex queries that require multiple joins and process large volumes of data. Modern database systems use heuristic methods together with superior optimization techniques to develop their database management systems.

Cost-Based Optimization

One of the most popular query optimization techniques in contemporary relational database systems is cost-based optimization (CBO). In contrast to heuristic optimization, cost-based optimization looks at a variety of execution strategies and selects the one with the lowest predicted execution cost [15]. Analyzers and preprocessors: The optimizer also creates multiple execution strategies based on varying access methods, join orders, indexing techniques, and finally lists the execution operators. The cost of each plan is computed using mathematical models that estimate resource consumption in terms of Disk I/O activities, CPU and memory consumption, and network connection overhead.

Cost estimation is largely dependent on statistical metadata maintained by the database system. Key statistics are the table cardinality, tuple distribution, histogram information, index selectivity and page access frequencies. These data help the optimizer

forecast execution costs and determine the size of intermediate outcomes.

Dynamic programming algorithms are frequently used in cost-based optimization to efficiently evaluate different join combinations. Modern DBMSs such as Oracle, MySQL, PostgreSQL, and SQL Server heavily depend on cost-based optimization for complex query processing.

Semantic Optimization

The semantic query optimization makes advantage of database semantic data, including integrity restrictions, functional dependencies, relationships and domain-specific knowledge, to enhance query performance optimization [16]. The semantic optimization process uses the logical characteristics of the data to identify which computations are unnecessary, rather than relying solely on statistical cost models.

The execution plans for queries become easier to handle because primary key and foreign key integrity constraints, uniqueness constraints and referential integrity rules exist. The foreign key relationship guarantees matching tuples, which allows safe removal of unnecessary join operations.

Materialized views are essential elements in semantic optimization. A materialized view stores precomputed query results that can be reused for future queries. The optimizer accesses data from the materialized view instead of re-executing complex aggregation and join calculations, resulting in faster execution times.

Semantic optimization also includes predicate simplification, elimination of redundant conditions, and query containment analysis. These techniques reduce processing complexity while enhancing execution performance.

Join Optimization

The execution of join operations in relational database systems requires more computational resources than any other database operation. The process of query execution needs efficient join optimization because it is vital for achieving high-performance results. The process of join optimization selects efficient join algorithms and determines the best sequence for executing joins.

Several join algorithms are commonly used in DBMSs:

Nested Loop Join: The nested loop join compares each tuple in one relation with every tuple in the other relation. The method demonstrates easy implementation, yet it becomes inefficient for large datasets because it requires extensive computational complexity.

Sort-Merge Join: The sort-merge join operation first requires both relations to be sorted on the join attributes, then merges matching tuples. The algorithm

achieves optimal performance when the input data is already sorted or when processing large data sets.

Hash Join: The hash joins method first partitions relations into hash buckets using hash functions and then performs efficient tuple matching. The method proves most effective when handling equi-join operations alongside large-dataset processing.

In addition to selecting join algorithms, query optimizers must determine the best join sequence. Different join orders produce different intermediate result sizes, which leads to execution performance being affected by join ordering.

Common join ordering strategies include:

- Left-deep join trees
- Right-deep join trees
- Bushy join trees
- Greedy algorithms
- Dynamic programming approaches

Efficient join optimization reduces execution time, memory consumption, and disk I/O operations significantly.

Indexing and Query Rewriting Techniques

The indexes are special data structures that help to speed up data retrieval without the need for a complete table scan. One of the most effective database optimization techniques used in relational databases is indexing to improve performance.

B-Tree Indexes: B-tree indexes function as balanced tree structures that enable users to perform range queries and ordered search operations. The system enables users to perform insertion, deletion and lookup operations with high efficiency.

Hash Indexes: Hash indexes bind keys to hash buckets, enabling exact-match searches. They work quite well for equality comparisons, but are unsuitable for range queries.

Bitmap Indexes: Bitmap indexes perform well for low-cardinality features commonly used in data warehouses and analytical workloads.

Clustered and Non-Clustered Indexes: Non-clustered indexes preserve distinct index structures that relate to data locations, while clustered indexes physically store data in sorted order.

Besides indexing, there are query rewriting techniques to transform SQL queries into more efficient ones. Examples of common query rewriting techniques are:

- Flattening nested subqueries
- Predicate simplification
- Eliminating redundant joins
- Common subexpression elimination
- Reordering selection conditions

The advantage of efficient indexing, plus query rewriting, is that it significantly enhances the performance of database systems, particularly in

enterprise applications with very large volumes of data.

Artificial Intelligence and Machine Learning-Based Optimization

In contemporary database systems, query optimization is boosted by the use of modern databases, which also leverage AI and ML methods, such as enhanced cost estimation, cardinality estimation and plan selection to optimize queries [17][18]. Unlike traditional methods, AI-driven techniques can adapt to fluctuating workloads and environments by leveraging query history. Neural networks are used to enhance cardinality estimation, and reinforcement learning is used to refine join-ordering strategies. AI-powered systems can also aid in automatic indexing, workload forecasting, database memory management, and self-tuning databases. While these methods enhance adaptability and performance, their practical application faces challenges such as complex training, high computational load, and the need for substantial data volumes.

Applications of Query Optimization Techniques

Applications of Query Optimization Techniques

Enterprise Database Systems: Many enterprise applications like banking, insurance, healthcare, retail and education systems use query optimization techniques extensively. Optimization leads to more efficient transactions, customer records and big data retrieval.

Data Warehousing and Business Intelligence: Analytical queries can be used in data warehouses and business intelligence platforms and can be optimized using optimization techniques[19]. Use of techniques like indexing, materialized views, and join optimization enhances reporting and decision-making.

Cloud and Distributed Databases: Query optimization is used in cloud-based and distributed database systems to minimize communication overhead, distribute the workload and enhance scalability across multiple servers and locations [20][21].

Big Data Analytics: Query optimization is the basis of the modern big data platforms that work efficiently with large data sets [22][23]. In systems like Hadoop or Spark, parallel processing and distributed query execution can aid in boosting processing speed.

Real-Time Processing Systems: Real-time applications such as stock trading, IoT monitoring, and online recommendation systems require optimized queries to ensure low latency and fast response times.

Challenges in Query Optimization Techniques

Cardinality Estimation Errors: Incorrect estimation of intermediate result sizes can lead to poor execution plan selection and inefficient query processing.

Dynamic Workloads: Modern applications experience continuously changing workloads, making static optimization techniques less effective.

Heterogeneous Data Sources: Database systems often integrate relational, NoSQL, graph, and streaming databases, increasing optimization complexity.

Large Search Space: Complex queries involving multiple joins generate a huge number of possible execution plans, increasing optimization overhead.

Resource Management: Efficient utilization of CPU, memory, storage, and network bandwidth remains a major challenge in large-scale database environments.

Emerging Trends in Query Optimization Techniques

AI-Driven Query Optimization: AI and ML are increasingly used for cardinality estimation, automatic indexing, and intelligent execution plan generation.

Autonomous Database Systems: Modern databases are evolving into self-driving systems capable of automatic tuning, optimization, and workload management.

Adaptive Query Processing: Adaptive optimization techniques dynamically modify execution plans during runtime based on changing workload conditions.

Big Data Query Optimization: Optimization strategies for big data systems focus on distributed processing, parallel execution, and in-memory analytics. A digital twin is considered as a live view of a real-world system that monitors the state of its entities. Deeply, it is an environment that consists of a virtual and a physical machine. Each machine (model) is represented as a simulation, a mirror, or a twin of the other. So, the digital twin can list the life cycle of the physical entity which can be a human, an object, or a process [68]. Each digital twin is connected to its counterpart by a unique key, therefore a relationship between two entities can be established [48]. A digital twin is a partition of a Cyber-Physical System (CPS), which is a set of physical systems connected to virtual cyberspace through the network [11, 49]. The communication between a physical entity and its digital twin can be represented directly by physical connections or indirectly via a cloud system. Also, it can be a seamless connection and continuous data exchange [26, 88]

Literature Review

The existing researches are presented that deals with the state of the art in query optimization in DBMSs, with emphasis on adaptive optimization, machine learning and evolutionary algorithms, distributed query processing, workload optimization and autonomous performance tuning, in order to achieve efficient database management.

Tian (2026) discusses the evolution and present status of query optimization in production systems and some of the problems being faced in production systems. Then this paper focuses on three trends that are on the rise within the industry that will have the potential to solve these problems. One, a new feedback loop is being employed between query optimization

and query execution, which is tighter to enhance query performance robustness. Second, a convergence of query optimization and workload optimization is expanding the scope of optimization from a single query to workloads. Third and most radically, the industry is shifting from monolithic architectures to composable architectures that promote agility and collaboration across engines [22].

Okoronkwo (2025) This topic's objectives are to examine and evaluate these three optimization levers' benefits and drawbacks, as well as their effects on distributed architectures like Hadoop, Spark, and Cloud SQL engines, as well as relational databases. The systematic review process was adhered to using the IEEE Xplore, ACM Digital Library, Springer, and Scopus databases. The Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework was followed to ensure that the procedure was transparent and repeatable. Following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) methodology, the procedure was carried out in a transparent and repeatable manner. Seven studies were eventually kept for synthesis out of the 5,908 records that were found, according to the study's inclusion and exclusion criteria. The data were systematically extracted, and notes were taken along with the writer, year, context of the database, optimization technique, and the main results of the research [23].

Hettiarachchi and Yapa (2025) It deals with two important agents in SQL QO, the Join Order Agent (JOA) and Index Selection Agent (ISA). The JOA seeks to determine the optimal join order among many tables to minimize the number of intermediate results and query execution time. The ISA is used to determine what indices are best suited to provide fast retrieval for different types of databases and queries. Survey the approaches that already exist for both agents, both heuristic-based and cost-based approaches, and ML approaches. Additionally, the study points out difficulties encountered in these aspects, including the scalability of current solutions, the ability to adapt dynamically to varying workloads and the incorporation of several optimization techniques [24].

Soepardi (2024) offers a detailed evaluation of several SQL query optimization techniques, focusing on evolutionary algorithms such as Genetic Programming (GP) and Genetic Algorithm (GA). While GP has been shown to be effective in producing innovative and unique query optimization strategies and solutions specific to the problem, the GA approach relies on mechanisms of selection, crossover and mutation to generate increasingly efficient query plans from one generation to next. While crossover strategies aim to combine the best features of several query plans, mutation is able to find other answers that may not be considered by conventional approaches [25].

Abhayanand and Rahman (2024). In a paper, a detailed survey of current methodologies, techniques and developments in query optimization for DRDBs is presented. Examine a variety of methods, including as adaptive optimization techniques, distributed query processing, cost-based optimization, and parallel query execution. They also talk about new developments, such as the integration of cloud computing technologies and optimization aided by machine learning. They want to find possibilities, gaps, and potential paths for improving query optimization in distributed relational databases through this research [26].

Huang et al. (2023) concentrate on database system performance optimization and examine important studies on prediction, diagnosis, and tweaking. Describe the strategies, benefits, and drawbacks of a number of suggested systems for both single and concurrent queries. Classify the methods by input data type, such as time metrics, logs, or monitoring metrics, and evaluate their diagnostic capabilities. Concentrate on the methods that operation and maintenance staff frequently use for tuning, such as storage management, elastic resources, knob tuning, index selection, view materialization, and SQL antipattern detection [27].

Table I shows the major research works on query optimization; focus, findings, challenges, limitations and future directions of research in modern distributed environments and database management systems.

Table 1 Summary of Literature Review on Query Optimization Techniques in Database Management Systems

Author(s)	Study on	Key Findings	Challenges	Limitations	Future Work
Tian (2026)	Review of historical and modern query optimization techniques in production database systems	Identified three major trends in query optimization: tighter feedback loops between optimization and execution, workload-level optimization, and composable architectures for cross-engine collaboration.	Maintaining robust query performance under dynamic workloads and heterogeneous systems.	Focuses mainly on industrial trends and conceptual directions rather than empirical benchmarking.	Development of adaptive and collaborative optimization frameworks across multiple database engines.
Okoronkwo (2025)	Systematic review of query optimization techniques in relational and distributed database	Evaluated strengths and weaknesses of cost-based, heuristic, and adaptive optimization techniques across Hadoop, Spark, and Cloud SQL systems using the PRISMA framework.	Managing optimization in distributed environments with scalability and transparency requirements.	Only 7 studies were retained after screening, limiting the breadth of synthesis.	More large-scale comparative studies and integration of optimization strategies in cloud-native DBMSs.

	systems				
Hettiarachchi and Yapa (2025)	Join Order Agent (JOA) and Index Selection Agent (ISA) SQL query optimization	Reviewed heuristic, cost-based, and machine learning methods for join ordering and index selection. Highlighted the importance of dynamic adaptation and integrated optimization.	Scalability of optimization methods and adapting to continuously changing workloads.	Existing techniques often optimize isolated components rather than integrated optimization pipelines.	Research on hybrid intelligent agents combining ML, cost-based, and adaptive optimization strategies.
Soepardi (2024)	SQL query optimization with evolutionary algorithms	Demonstrated that Genetic Programming (GP) and Genetic Algorithms (GA) can generate efficient and innovative query execution plans through crossover and mutation operations.	High computational complexity and convergence time in evolutionary optimization methods.	Evolutionary methods may require extensive training iterations and system resources.	Exploration of hybrid evolutionary and machine learning-based optimization techniques for real-time query tuning.
Abhayanand and Rahman (2024)	Distributed Relational Database Systems (DRDBs) query optimization	researched parallel execution techniques, adaptive optimization, distributed query processing, and cost-based optimization. Discussed cloud integration and ML-assisted optimization trends.	Efficient query execution across geographically distributed databases and cloud environments.	Review-oriented study with limited experimental validation of discussed methods.	Investigation of AI-driven distributed optimization and scalable cloud-native query processing frameworks.
Huang et al. (2023)	Database performance optimization including prediction, diagnosis, and tuning	Categorized optimization approaches into prediction, diagnosis, and tuning. Reviewed techniques such as knob tuning, index selection, SQL anti-pattern detection, and elastic resource management.	Accurate diagnosis and tuning for concurrent queries and dynamic database workloads.	Some reviewed methods are system-specific and difficult to generalize across platforms.	Development of autonomous self-tuning and intelligent database optimization systems using AI techniques.

Conclusion and Future Work

Query optimizing databases is an essential part of Database Management Systems (DBMS) that significantly impacts query performance, system scalability, and resource utilization. In this paper, the main query optimization techniques, such as heuristic-based optimization, cost-based optimization, semantic optimization, join optimization, indexing methods and query rewriting strategies, were reviewed. The study also examined modern AI- and Machine Learning-based optimizations designed to enhance query performance adaptively, as well as cardinality estimation, auto-tuning, and workload management. In enterprise databases, cloud computing, distributed systems, data warehouses, and big data analytics systems where high-performance data processing is crucial, efficient query optimization is a critical aspect. Even with significant progress, there are still open issues such as poor cost estimation, dynamic workloads, large search spaces for execution plans, dealing with heterogeneous data integration and efficient resource management in distributed environments. A number of these limitations are anticipated to be resolved by emerging technologies like autonomous databases, adaptive query processing, and intelligent optimization frameworks. Future research directions include the development of scalable AI-driven optimization models, hybrid optimization methods, real-time adaptive execution methods, and energy-efficient query processing methods. Moreover, the marriage between machine learning and classical optimization methods could be used to improve the robustness, flexibility and overall

performance of future database systems that handle large and dynamic datasets.

References

- [1] M. Chanda, "A Low-Cost System for Acquiring Login/Logout Data for On-Ground Racks of in-Flight Entertainment Systems," California State University, 2016.
- [2] A. K. Padhy, C. Medicherla, B. Vulugundam, C. Kulkarni, T. P. Patel, and S. Shivam, "Latency-Optimized Microservices Orchestration for Real-Time E-Commerce in Multi-Cloud Environments," in 2025 International Conference on Computer and Applications (ICCA), 2025, pp. 1–6. doi: 10.1109/ICCA66035.2025.11430930.
- [3] A. Parupalli, "Improving Cloud Resource Utilization and Cost Efficiency Through Data-Driven ML Algorithms," in 2025 Global Conference on Information Technology and Communication Networks (GITCON), 2025, pp. 1–8. doi: 10.1109/GITCON65266.2025.11376599.
- [4] N. Radhasharan, "Real-Time Edge-To-Cloud Intelligence Architecture For Autonomous Drilling Systems," J. Int. Cris. RISK Commun. Res., vol. 9, no. 1, pp. 90–102, 2026, doi: 10.63278/jicrcr.vi.3577.
- [5] H. B. Dama, "A Survey of MySQL Database Administration Techniques and Best Practices," ESP J. Eng. Technol. Adv, vol. 6, no. 1, pp. 89–98, 2026.
- [6] P. Naayin and S. Kamatal, "Automating Infrastructure Platforms with Cloud, Kubernetes, and Site Reliability Engineering," Int. J. Comput. Tech., vol. 8, no. 6, 2021.
- [7] C. Patel, "A Review of Multi-Channel CRM Strategies Using Big Data and Cloud Integration," Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol., vol. 8, no. 1, pp. 577–588, 2022.
- [8] H. Ravilla, "Predictive Analytics for Customer Churn in Salesforce Service Cloud," in Data Science and Big Data Analytics, D. Mishra, X. S. Yang, A. Unal, and D. S. Jat, Eds., in Learning and Analytics in Intelligent Systems, vol. 55. Cham: Springer, 2026, pp. 14–24. doi: 10.1007/978-3-032-05377-0_2.

- [9] V. K. Sharma, "Cloud Computing \& IoT: 5G Focused IoT with Cloud Solutions," *Int. J. Artif. Intell. Big Data, Cloud Comput. Data Sci.*, vol. 6, no. 3, pp. 21–25, Jul. 2025.
- [10] J. B. Mehta, "AI-Driven Compliance Validation System and Framework for Multi-Cloud Enterprises," 63/919,564, 2025
- [11] T. P. Patel and P. Agarwal, "DeepServe: Hierarchical Model Placement and Dynamic Batching for Cost-Efficient Multi-Tenant LLM Inference at Scale," in *SoutheastCon 2026*, 2026, pp. 1–6. doi: 10.1109/SoutheastCon 63549.2026.11476566.
- [12] B. Krishnan, S. Perla, S. Maddela, and R. Lingam, "Adaptive Multi-Cloud Infrastructure for CRM Analytics: Real-Time ML and Data Sync with LLMs," in *2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN)*, IEEE, Nov. 2025, pp. 1–8. doi: 10.1109/GWCN66157.2025.11448404.
- [13] M. R. C. Mukkolakkal, "IntelliStore: An Intelligent AI Agent Framework for Autonomous Storage and Database Optimization in Cloud-Native Microservices," *Int. J. Sci. Res. Mod. Technol.*, vol. 3, no. 12, pp. 243–250, 2024, doi: 10.38124/ijsrmt.v3i12.1024.
- [14] S. J. Kamatkar, A. Kamble, A. Viloría, L. Hernández-Fernández, and E. G. Cali, "Database performance tuning and query optimization," in *International Conference on Data Mining and Big Data*, 2018, pp. 3–11.
- [15] M. M. F. Fahima, A. H. S. Sreen, S. L. F. Ruksana, D. T. E. Weiheña, and M. H. M. Majid, "Machine learning for database management and query optimization," *Elem. J. Educ. Res.*, vol. 2, no. 1, pp. 96–108, 2024.
- [16] S. Kilaru, V. S. R. Yarlagadda, and M. B. Nagaraja, "AI-Augmented Block-Sketch Hybrid Compression: Enabling Semantic Query Processing on Compressed Data," in *2026 IEEE 16th Annual Computing and Communication Workshop and Conference (CCWC)*, 2026, pp. 966–972. doi: 10.1109/CCWC67433.2026.11393750.
- [17] A. Nerella, N. Kolli, and J. W. Sajja, "Building Secure AI Agents for Autonomous Data Access in Compliance/Regulatory-Critical Environments," *SSRN Electron. J.*, vol. 2024, no. 9, Sep. 2024, doi: 10.2139/ssrn.5528763.
- [18] A. Warriar, "Agentic AI for Natural Language Query Interface in Intelligent Customer Success Management-Conversational Analytics and Automated Reporting," *J. Artif. Intell. Cloud Comput.*, vol. 450, pp. 1–6, Aug. 2025, doi: 10.47363/JAICC/2025(4)450.
- [19] T. Shah, "Cloud-Based Data Warehousing for Marketing Agility: Lessons from FinTech Migrations to Snowflake and AWS," *Int. J. Adv. Res. Sci. Commun. Technol.*, vol. 4, no. 4, pp. 642–652, 2024, doi: 10.48175/IJARST-16000B.
- [20] S. Singamsetty, "An Intelligent Framework for Secure and Fair Cloud Resource Distribution," in *2025 7th International Conference on Innovative Data Communication Technologies and Application (ICIDCA)*, IEEE, Oct. 2025, pp. 686–690. doi: 10.1109/ICIDCA66325.2025.11280502.
- [21] V. Sanikal, "AI-Enhanced Network Security Framework for Cloud-Edge Integrated Environments," in *2026 Innovations in Machine, Engineering, and Digital Conference (IMED)*, 2026, pp. 1–6. doi: 10.1109/IMED68921.2026.11484417.
- [22] Y. Tian, "Query Optimization in the Wild: Realities and Trends," arXiv, 2026
- [23] M. Okoronkwo, "Next-Generation Query Optimization: A Systematic Review of Cost Models, Equivalence Rules, and Materialized Views," *Int. J. Comput. Sci. Math. Theory*, vol. 11, no. 9, pp. 56–69, 2025, doi: 10.56201/ijsrmt.vol.11.no9.2025.pg56.69.
- [24] N. Hettiarachchi and P. Yapa, "Advancements in SQL Query Optimization: A Review of Join Order and Index Selection," in *2025 5th International Conference on Machine Learning and Intelligent Systems Engineering (MLISE)*, 2025, pp. 31–39. doi: 10.1109/MLISE66443.2025.11100192.
- [25] F. D. T. Soepardi, "Evolutionary Algorithms for SQL Query Optimization: A Systematic Literature Review," 2024. doi: 10.13140/RG.2.2.12807.20646.
- [26] Abhayanand and D. M. M. Rahman, "Enhancing Query Optimization in Distributed Relational Databases: A Comprehensive Review," *Int. J. Nov. Res. Dev.*, vol. 9, no. 3, pp. 590–602, 2024.
- [27] S. Huang, Y. Qin, X. Zhang, Y. Tu, Z. Li, and B. Cui, "Survey on performance optimization for database systems," 2023. doi: 10.1007/s11432-021-3578-6.