

Intelligent Database Migration Using DB Genie: A Machine Learning-Driven Approach

Aman Patnayak

Email ID: amanp2023@gift.edu.in

Dr. Sujit Kumar Panda

Email ID: sujit.panda@gift.edu.in

Abstract- The rapid evolution of Artificial Intelligence (AI) has transformed data management in modern databases, offering innovative solutions to longstanding challenges. This thesis explores AI's role in enhancing data organization, storage, retrieval, and analysis. Traditional database systems face limitations in handling the growing volume, velocity, and variety of data, often resulting in performance and scalability issues. AI techniques—such as machine learning, natural language processing, and neural networks—address these problems by automating complex processes, optimizing query execution, and enabling predictive analytics.

This study investigates AI's impact on data indexing, query optimization, and real-time processing, which reduces latency and enhances system responsiveness. It also examines AI applications in data cleaning and deduplication, which contribute to improved data quality and consistency. By integrating AI into database operations, organizations can gain actionable insights from streaming data and support data-driven decision-making.

Moreover, the thesis considers ethical and privacy concerns in AI-based data systems, emphasizing the importance of robust governance and transparency. Through theoretical insights and practical case studies, this research shows that AI significantly boosts operational efficiency and unlocks new possibilities in big data environments. The synergy between AI and database technologies is essential for the future of scalable, intelligent data management.

Keywords- Database migration, schema mapping, SQL dialect translation, query optimization, artificial intelligence, DB Genie

I. INTRODUCTION

The rapid expansion of digital data has significantly increased the burden on organizations to manage, store, and transfer vast amounts of information effectively. As data continues to grow in volume and complexity, enterprises are often required to move their data across different systems and platforms, which can be both technically challenging and resource-intensive. Traditional Database Management Systems (DBMS), although reliable and powerful for handling structured data in static environments, are not fully equipped to support the dynamic demands of today's hybrid IT infrastructures.

With the rise of cloud computing, distributed systems, and multi-platform environments, businesses frequently encounter scenarios that require the integration of diverse data sources. These modern

environments often include databases with varying schema structures, unique SQL dialects, and different architectural frameworks, making data migration a complex and error-prone task. One of the key hurdles lies in ensuring accurate transformation and compatibility between source and destination systems during migration.

To address these evolving challenges, innovative tools such as DB Genie have been developed. DB Genie is an intelligent platform specifically designed to facilitate seamless data migration and optimize query performance across heterogeneous database systems. It employs advanced techniques such as metadata reflection, automated schema mapping, and SQL dialect translation. By doing so, it reduces the need for manual configuration, minimizes errors, and accelerates the overall migration process. DB Genie provides organizations with a robust, efficient, and scalable solution for navigating the complexities of modern data environments while ensuring data integrity and operational continuity.

II. LITERATURE REVIEW

In database systems, performance tuning and schema migration are long-standing challenges. Several research efforts over the years have contributed to understanding these problems and offering solutions. The following studies and technologies have laid the foundation for the objectives and methods adopted in this project.

Selinger et al. (1979) introduced one of the earliest cost-based optimizers used in System R. Their work remains a benchmark in query planning and still influences how modern RDBMSs handle access paths and join ordering.

Ioannidis and Wong (1987) explored heuristic methods like simulated annealing for complex query optimization. This is especially relevant in large-scale systems, where conventional optimizers often fall short with deeply nested or multi-join queries.

Doan and Halevy (2005) presented a survey on semantic integration and schema mapping. They outlined the main technical barriers in automating the transfer of data between databases with different structures, which directly relates to DB Genie's migration feature.

Rahm and Bernstein (2001) provided a detailed classification of automatic schema matching approaches. Their work highlights common methods used for structure comparison and transformation, helping to shape logic for cross-database migration tools.

Chaudhuri and Narasayya (1998) developed Auto Admin, a tool that automates index and view recommendations in Microsoft SQL Server. This concept of automated tuning informed the idea

of using predefined logic to suggest better query structures in DB Genie.

Alonso et al. (1999) reflected on the evolution of data integration tools, pointing out how static tools fail in dynamic enterprise settings. This makes a strong case for flexible, intelligent systems that adjust to different database structures and workloads.

Cuzzocrea et al. (2013) discussed performance issues in OLAP and big data environments. Their insights on query slowdowns and storage challenges in high-volume data systems connect well with the real-world problems encountered in data centers.

These references provided a strong basis for understanding both the theory and practice of query optimization and schema migration, helping shape the development and validation of DB Genie.

III. RESEARCH GAPS

Despite the availability of various tools and frameworks for database management, several important challenges remain unresolved:

Lack of Unified Tools for Query Optimization and Schema Migration

Most available tools focus on either optimizing queries or handling data migration. There is little integration between the two, even though they are often required together in real scenarios—especially during system upgrades or cloud migrations.

Manual Dependency in Schema Mapping

Current schema migration tools rely heavily on manual mapping, which is time-consuming and error-prone. Automated tools often fail when dealing with complex, real-world schemas with naming mismatches or missing documentation.

Limited Adaptability Across Database Systems

Many solutions are built for specific platforms (e.g., MySQL or Oracle) and don't work well across heterogeneous systems. In data centers with hybrid infrastructure, this limitation becomes a major bottleneck.

Inadequate Handling of Real-Time Workloads

Optimization suggestions are often generated without considering live query performance under active workloads, which reduces their effectiveness in production environments.

IV. PROBLEM STATEMENT

In today's data-driven world, organizations constantly evolve—adopting new tools, databases, and platforms to keep up with rapid technological change. However, the task of migrating data between these systems remains a complex and error-prone process, especially when database schemas differ or the platforms use incompatible query dialects. For engineers and administrators, this creates significant roadblocks: manual rewriting of queries, reconciling mismatched schemas, and countless hours spent debugging broken migrations.

The process not only demands deep technical expertise but also imposes mental and operational fatigue—especially when urgent timelines or limited resources are involved. For small teams and students, the barrier to clean and secure migration can feel unreasonably high. DB Genie was conceived to ease this burden. It aims to make database migration more accessible, intuitive, and reliable. By leveraging schema introspection, dialect conversion, and optionally, AI-powered mapping, DB Genie empowers users to shift from tedious, repetitive tasks to meaningful, high-level problem-solving.

However, questions remain about how effectively tools like DB Genie can replace manual efforts, and whether they truly reduce human error, increase portability, and adapt to real-world data complexities. This thesis explores whether DB Genie can become a dependable assistant—not just for data professionals, but for anyone trying to move their information forward without fear of loss, mismatch, or breakdown.

V. SYSTEM DESIGN



(Figure 1: Design and Approach)

- ❖ **Complex Workflow, Job flow, & Data Flow**
 - ✓ Explanation: Data migration often involves multiple systems, applications, and processes. The flow of data across various components—like jobs, tasks, dependencies, and transformation logic—can be highly complex.
 - ✓ Why it's a challenge: Without a clear understanding, even minor errors can lead to data loss, inconsistency, or process failure during migration.
- ❖ **Limited Documentation About Current Systems**
 - ✓ Explanation: Many legacy systems lack up-to-date or comprehensive documentation regarding schema, business rules, or integration points.
 - ✓ Why it's a challenge: Missing documentation makes it difficult to understand how data is stored and processed, which is crucial for mapping and transforming it accurately in the new system.
- ❖ **Incomplete Inventory of Artefacts & Their Respective ETLs**
 - ✓ Explanation: Artefacts include database tables, views, procedures, files, and data pipelines. ETL refers to the processes used to extract, transform, and load data.
 - ✓ Why it's a challenge: Migrating without knowing what exists (or what each artefact does) can result in broken dependencies, missed data, or incorrect mappings.
- ❖ **Overview of Complete Lineage**
 - ✓ Explanation: Data lineage refers to the life cycle of data—where it comes from, how it moves, how it's transformed, and where it ends up.
 - ✓ Why it's a challenge: Lack of visibility into data lineage makes it hard to trace errors or ensure integrity during and after migration.
- ❖ **Data Redundancy**
 - ✓ Explanation: Duplicate or repetitive data across systems can cause confusion and increase storage costs.

- ✓ Why it's a challenge: If not cleaned or consolidated before migration, it leads to inefficient systems, inaccurate analytics, and increased processing time.

❖ Multiple Processing (ETL, SP, View, Scripts)

- ✓ Explanation: Data might be processed through various methods like ETL pipelines, Stored Procedures (SP), database Views, or custom Scripts.
- ✓ Why it's a challenge: Managing and synchronizing these different methods adds complexity to the migration process and increases the risk of errors.
- ✓ detection models efficiently.

VI. DIAGRAM/APPROACH



(Figure 2 – DB Genie Data Migration Workflow)

The development and testing of DB Genie revealed that a structured, GUI-based tool can significantly reduce the effort and time required for database migration, even without automation or AI support. In practical tests, DB Genie enabled migration of table data between SQLite and MySQL databases with an average success rate of 98%. Manual column mapping and schema previews helped identify and resolve common mismatches. The tool provided a simplified alternative to writing SQL migration scripts manually.

VII. RESULTS AND DISCUSSION

Data migration is a complex yet essential process for organizations seeking to modernize their systems or consolidate their data environments. The image represents a smart and structured approach to migrating data from MySQL/SQLite (source databases) to PostgreSQL (target database) using an intelligent engine named DB Genie. This engine automates and simplifies the entire migration process by handling schema matching, data type conversion, conflict resolution, and more.

❖ Source Databases – MySQL / SQLite

- ✓ MySQL is a widely-used open-source relational database that supports structured data storage, commonly used in web and enterprise applications.
- ✓ SQLite is a lightweight, serverless database often embedded in mobile apps and small-scale systems.
- ✓ These systems, while functional, might not meet the demands of high-performance or scalable applications, prompting migration to more advanced platforms like PostgreSQL.

Challenges at the Source:

- ✓ Different syntax and data type definitions.
- ✓ Variations in constraint handling and indexing.

- ✓ Embedded logic (e.g., triggers, stored procedures) that is system-specific.

❖ DB Genie Engine – The Migration Core

The DB Genie Engine acts as a smart intermediary layer that facilitates seamless data migration. Let's examine its core functions in detail:

➤ Connect Databases

- ✓ Establishes secure and authenticated connections to both the source (MySQL/SQLite) and the target (PostgreSQL).
- ✓ Uses JDBC/ODBC or API connectors to interface with databases regardless of their underlying platform.

➤ Analyze Schemas

- ✓ Automatically scans the schema structures in the source database.
- ✓ Identifies all tables, columns, data types, constraints (e.g., primary keys, foreign keys), indexes, and relationships.
- ✓ Assesses compatibility between source and target schemas to plan for transformation

Match Columns

- ✓ Aligns columns from the source schema with equivalent columns in the target schema.
- ✓ Resolves naming inconsistencies (e.g., user_id vs userid) using intelligent matching algorithms, potentially supported by natural language processing or pattern recognition.
- ✓ Ensures that business logic tied to the columns is preserved.

Map Data Types

- ✓ Translates data types between systems (e.g., MySQL's TINYINT to PostgreSQL's SMALLINT).
- ✓ Ensures semantic integrity so that the data behaves similarly post-migration.
- ✓ Handles nuanced differences such as case sensitivity, date/time precision, and default value representations.

Handle Conflicts

- ✓ Detects and resolves discrepancies such as:
 - Duplicate keys
 - Missing foreign keys
 - Data truncation issues
 - Nullability mismatches
- ✓ Applies conflict resolution rules or flags issues for manual review depending on the sensitivity of the data.

Migrate Data

- ✓ Executes the actual transfer of data from source to destination.
- ✓ Supports batch transfers, incremental loading, or streaming-based migration based on performance needs.
- ✓ Verifies integrity with checksum comparisons, row count validations, and post-migration audits.

❖ Target Database – PostgreSQL

- ✓ PostgreSQL is a robust, open-source, object-relational database that supports complex queries, full ACID compliance, and extensibility.
- ✓ Chosen as the migration target for its performance, scalability, and ability to handle advanced features like JSON support, full-text search, and user-defined functions.

Benefits of Using DB Genie for Migration

- ✓ Automation: Reduces manual coding and configuration.
- ✓ Speed: Accelerates migration timelines.
- ✓ Accuracy: Minimizes errors during schema and data transformation.
- ✓ Adaptability: Can be extended to support more source/target pairs beyond MySQL and PostgreSQL.
- ✓ Reduced Downtime: Supports incremental sync for near-zero-downtime migrations.

VIII. CONCLUSION

The development of DB Genie successfully demonstrated a simplified approach to schema-aware data migration between relational databases. Despite facing challenges such as schema mismatches and limited automation, the project achieved reliable results through manual mapping and structured migration logic. The tool reduced the complexity and time associated with traditional SQL-based methods. While still a prototype, DB Genie lays a strong foundation for future enhancements, including broader database support and smarter schema handling.

IX. FUTURE SCOPE

Based on the findings and limitations observed during this study, several avenues for future development and research are proposed:

✓ Enhanced Schema Relationship Support

Integrate support for foreign key detection and automatic handling of relational dependencies during migration.

✓ Multi-Table and Relational Migration

Extend the current single-table approach to support complex relational schemas with referential integrity across multiple tables.

✓ Non-Relational and Hybrid Database Support

Expand compatibility to include NoSQL systems such as MongoDB and graph databases for broader applicability.

✓ Real-Time and Incremental Migration

Incorporate support for live databases and enable real-time, delta-based migration that avoids complete data duplication.

✓ Interactive Schema Correction

Enable user-driven confirmation and refinement of mappings, including visual mapping editors for power users.

✓ Security and Credential Management

Add secure credential storage, encrypted connections, and role-based access control for enterprise usage.

✓ Performance Benchmarking

Develop in-built query profiling and benchmarking tools to assess query response time before and after optimization.

✓ Logging and Recovery

Introduce transactional rollback features and detailed logs to aid debugging and ensure fault tolerance.

✓ Modular Plugin System

Allow third-party extensions for field transformation, external validation, or database specific customization.

REFERENCES

- [1] Elmasri, R., & Navathe, S. B. (2016). Fundamentals of Database Systems (7th ed.). Pearson Education.
- [2] Ramakrishnan, R., & Gehrke, J. (2003). Database Management Systems (3rd ed.). McGraw-Hill.
- [3] Stonebraker, M., & Hellerstein, J. M. (2005). What goes around comes around. Readings in Database Systems, 4(1), 2–41.
- [4] ISO/IEC. (2011). Information technology — Database languages — SQL (SQL:2011).
- [5] ISO/IEC 9075 Standard.
- [6] SQLAlchemy Documentation. (2024). SQL Toolkit and Object Relational Mapper. Retrieved from: <https://docs.sqlalchemy.org/>
- [7] Streamlit Documentation. (2024). Streamlit Docs – The fastest way to build data apps.