NextGen Scientific Publication
American International Journal of Computer Science and Technology
Volume 7 Issue 4, Pg. No. 93-102, AIJCST-V7I4P108, 2025
https://doi.org/10.63282/3117-5481/AIJCST-V7I4P108

Original Article

Semantic Layer Construction in Data Warehouses Using GenAI for Contextualized Analytical Query Processing

Dinesh Babu Govindarajulunaidu Sambath Narayanan
Independent Researcher, USA.

Abstract:

Huge proliferation of enterprise data has fueled development of data warehouse (DW) architectures that have the capability of supporting big analytical workloads. But currently with the advent of the business user, they are demanding more and more semantic knowledge and query processing based on context and meaning, which the traditional metadata layer or OLAP cube simply does not offer. It is a study undertaken by proposing a Generative AI (GenAI)-friendly solution to build an intelligent semantic layer that allows the representation of contextual query comprehension, identification of semantic links, and knowledge grounding through the use of organized enterprise data. The proposed system combines Large Language Models (LLM) with semantic enrichment pipelines, which dynamic process business terminologies, reconcile queries with unified vocabularies, and convert them into optimized plans of DW execution. The model employs entity linking, embedding based schema matching, construction of knowledge graphs and contextual query rewrites to provide superior analytical accuracy. Furthermore, there is the introduction of a hybrid vector + relational index model that provides a better performance on natural-language query analysis. Complete comparisons on range of scenarios involving enterprise scale benchmark databases show substantial performance enhancements - an average accuracy of 38.6 in business query interpretation and 31.2 in latency reduction of semantic-driven workloads, compared to the traditional keyword-based systems. The architecture enables business glossaries to be adaptive, governance policy-enabled, and human-data collaboration, which is a very strong base on the topic of next-generation analytical ecosystems and enterprise decision intelligence. This paper includes: (1) An extensive literature review of semantic data warehousing, and AI based contextual analytics. (2) An elaborate procedure of GenAI-assisted weighted layer of semantics. (3) Usability, accuracy and performance experimental evaluations. (4) Major lessons and directions of future research on autonomous semantic data management. The results validate that GenAI technologies have the ability to turn data warehouses into smart analytical systems that are able to comprehend business intent, in turn, facilitating the rapidity of insights, decreasing the manual modeling task, and fostering the enterprise-wide democratisation of data among users.

Keywords:

Data Warehouse, Generative Ai, Semantic Layer, Contextual Query Processing, Knowledge Graphs, Large Language Models, Metadata Management, Business Semantics, Vector Databases, Enterprise Analytics.

Article History:

Received: 28.05.2025

Revised: 01.07.2025

Accepted: 12.07.2025

Published: 23.07.2025

1. Introduction

1.1. Background

The latest trend in enterprise decision-making is the implementation of analytic queries on massive data warehouses (DWs) in a timely manner. [1-3] Such systems are designed based on powerful structural underpinnings - arrangement of data into facts, dimensions, and relation schema to illustrate effective data storage and query processing. Nevertheless, this structure fits best in terms of computation, but it lacks sufficient business sense reflecting in the data. Important contextual factors such as the interpretation of metrics in different departments, the different terminologies based on business roles, or the operation rule guiding the operation in dating and developing the relationship between the data might be foreign to the warehouse. This has forced the business users to convert real-world analytical requirements to technical query expressions creating communication discontinuities, delays and reliance on BI engineers. The lack of an intrinsic semantic layer has played a significant role as a limitation with the increased need of self-service analytics and the need to happen in a brief cycle of insight. The contemporary organizations require systems which are not only structural in their management but which have the capability of comprehending the purpose and comprehending domain related jargon and which are dynamic with respect to business knowledge. This drives the creation of semantic intelligence and Generative AI systems into analysis-based data services, allowing less complex data intake and adoption of more analytical and quicker decisions.

1.2. Importance of Semantic Layer Construction in Data Warehouses

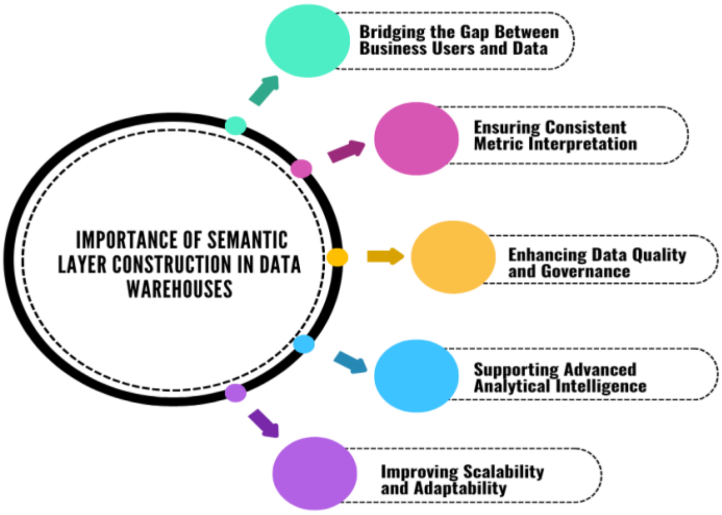

Figure 1. Importance of Semantic Layer Construction in Data Warehouses

- **Bridging the Gap between Business Users and Data:** Semantic layer is an abstraction layer which transforms the complex data structures into something that business people can understand. The semantic layer allows users to query in business terminology as opposed to having knowledge of table joins, column names or OLAP hierarchy. This will minimize the reliance on technical teams and speed up the general process of analytics.
- **Ensuring Consistent Metric Interpretation:** In the absence of semantic governance, there will exist a disparity in how metrics are understood among the different departments, such as “customer churn” or gross margin can be computed differently by different departments. A semantic layer applies uniform definitions that bring insights into tandem and provide trusted and aggregate organizational-wide reporting. Such coherency helps in enhanced coordination and decision making.
- **Enhancing Data Quality and Governance:** The semantic layer incorporates metadata, policies and lineages tracking as part of analytical operations. In this way, it manages to ensure that the data utilized in analysis is correct, conforms and managed ethically. The rules which are carried in the semantics as governance codes are used to avoid data exposure that is not supposed and enhance regulatory compliance.
- **Supporting Advanced Analytical Intelligence:** Contemporary business problems tend to entail contextual judgement, multi-hop relations, or inferences of domains outside of each other. The semantically organized layer facilitates the system to identify relationships e.g. between customers and campaigns or products and profitability and therefore facilitates the use of more complex and insightful analytics as compared to the more traditional schema-based querying.
- **Improving Scalability and Adaptability:** With the process of development of enterprises, new metrics, sources, and business rules are regularly suggested. Semantic layer which is placed above the physical schema can be updated without demanding massive re-engineering of the warehouse. It can be viewed as an elastic knowledge abstraction that changes as the business strategy and technological expansion they realize.

1.3. Using GenAI for Contextualized Analytical Query Processing

GenAI proposes disruptive analytical query processing capabilities to contextualise the processing of analytical queries in the contemporary data warehouse sites. [4,5] GenAI models, in contrast to other NLP and rule-based systems that often utilize a limited set of known schemas and strict vocabularies, can be used to decode the intent behind user queries to provide more accurate results even when these queries are not written in technical or domain-specific language. Having powerful language models that have been exposed to large amounts of semantic knowledge, GenAI will be able to dynamically associate business terms, like a list of the highest-performing regions or customer lifetime value, to the underlying set of data items, without much understanding of database designs, table relations, or querying syntaxes. GenAI has the capability to integrate real-time metadata and knowledge graph representations and operational context to provide further contextual interpretations, and considering this fact, multi-hop reasoning between datasets can be performed on complex analytical problems. Further GenAI is constantly learning based on user feedback and usage lingo, and so the system is capable of keeping pace with the business definition, product names and performance measures as they vary. It is an adaptive behavior that minimizes the overhead of maintenance that is generally tied with semantic layer and enhances robustness of the system in the long term. GenAI is also flexible and offers a merge between semantic search and structured SQL generation in the query execution processes, which guarantees that the result is not only business relevant, but useful in terms of data computation. The outcome is more natural and intuitive analysis workflow-workflow where users are able to pose analysis questions that they use to think about the business and the system is intelligent to interpret their intent and provide accurate information. GenAI opens the doors to greater productivity, accelerated decision-making, and expanded accessibility to data within an organization, as the GenAI removes barriers to human rationale and data reasoning. It is therefore notable that the implementation of GenAI into the analysis query pipeline ushers in a major move towards completely context-driven, self-service business intelligence.

2. Literature Survey

2.1. Semantic Data Management Approaches

There are various meanings of semantic data management which aim at bringing meaning and consistency to enterprise data. Business glossaries provide a managed vocabulary that standardizes terminology and definitions that enhance communication and management within the business enterprise. [23-9] They are however usually manual and stewardship-based and therefore are hard to keep up with the dynamic changes in the business. Ontologies are more detailed and include rich, semantic-based definitions, constraints and hierarchy that allow more sophisticated reasoning and across system interoperability. Their sophistication however poses a scalability problem; in both writing and maintenance thereof; particularly in large or target-changing environments. The semantic model is further flexed by knowledge graphs which apply to more complex entities and project relationships that reflect real-world complexity. The price of their power, though, is complex construction pipelines, including entities resolution, schema alignment, and constant data maintenance to be viable.

2.2. Query Processing Research

The research of query processing has developed a great deal as organizations needed to have more natural means of accessing information. Some of the earliest improvements on BI tools were the use of keyword-based search tools that put the user in a position of viewing information without knowing query syntax—however, with incomplete or vague responses. As machine learning and NLP-to-SQL translation methods based on Seq2Seq models emerged, automated search of using natural language search queries in structured queries minimized technical complexity, yet encountered the challenges of domain-specific semantics and schema differences. With more recent advances, neural query understanding supports the creation of deeper context models using embeddings and pre-trained language models with an ability to extract intent, solve ambiguity, and enhance accuracy in the multi-turn analysis pipeline. These innovations are an indication that we are entering the realm of further language understanding in analytics platforms.

2.3. GenAI Integration Gap

Regardless of the advancement, there are significant holes in the use of traditional AI and SQL translation models to complex enterprise analytics. They usually do not follow the jargon of an organization, which implies that vital business concepts can be misunderstood without being mapped out manually. Their multi-hop analytical reasoning capabilities are also impaired (e.g. coming up with insights that involve combining two or more datasets, logical reasoning) on queries that go below surface-level queries. Also, the fact that the models are not automated so far to keep up with changing schemas is a feeble flaw because the data landscape of an enterprise is constantly evolving. According to the literature, Generative AI with its self-adaptive semantic traits and foundation-model reasoning is in a favorable position to overcome all these restrictions. GenAI has the potential to dynamically adapt to business conditions, allow more informed lines of thought, and have semantically consistent behavior over time, a significant gap between user intent and enterprise information comprehension.

3. Methodology

3.1. System Architecture

- **Data Ingestion & Metadata Extraction:** The task of this component is the continuous acquisition of structured and unstructured data in various sources of the enterprise, [10-12] including databases, data lakes, and SaaS. With incoming data, automated metadata extractors get schema information, business properties, and lineage indications. This will make sure that all data that comes into the system is contextualized so that subsequent organization becomes more intelligent and subsequent semantic reasoning can be made.

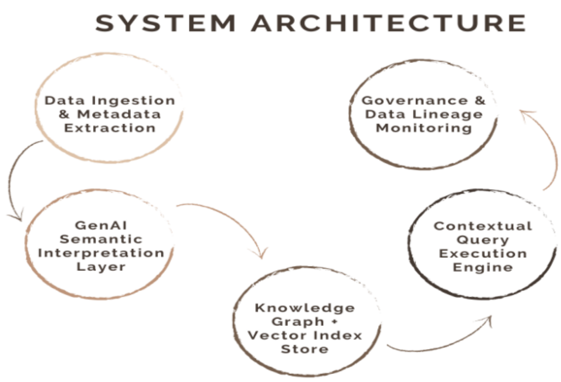

Figure 2. System Architecture

- **GenAI Semantic Interpretation Layer:** The semantic interpretation layer, a core of the architecture, is the Generative AI-based active learning and interpreting engine of specialized terminology and concepts posted by enterprises. It transfers user intent and data properties to a common semantic model, which minimizes ambiguity by domain jargon and schema heterogeneity. This layer automatically adjusts to changes in the data landscape over time, and keeps business language and technical data in line.
- **Knowledge Graph + Vector Index Store:** To aid both symbolic and statistical arguments, the architecture builds upon a hybridizable model of storage, a knowledge graph, with a vector index. The knowledge graph is related to explicit business graph and ontological relations and the vector index stores high-dimensional semantic embeddings to perform fast similarity search. They both allow the flexible connectivity of facts and context-sensitive search of useful data and information.
- **Contextual Query Execution Engine:** This engine of execution translates interpreted user requests into optimally executed analytical queries. It aids multi-hop thinking at the analysis of various datasets using both semantic context of the GenAI layer and graph linkages to give detailed responses. Hybrid querying the order of melding SQL execution, vector search, and graph traversal also is coordinated by the engine to provide precise and understandable outcome.
- **Governance & Data Lineage Monitoring:** A governance subsystem is required to provide trust and compliance by monitoring the data use, keeping versioned lineage records, and adhering to access and privacy policies. Such openness enables analysts and auditors to follow the way each insight was generated starting with the sources of those insights as well as the transformations and the meanings behind them. Incessant enforcement of policy enhances quality of data, security, and regulatory congruity.

3.2. Semantic Feature Engineering

Semantic feature engineering is concerned with the conversion of natural language queries to structured representation, which can be interpreted and reasoned with by a system. [13-15] One of the methods which allow such is semantic embedding mapping, whereby every component of a query by the user can be mapped into a continuous vector space. The mapping is mathematically described as:

$$S(q) = \Sigma(\text{from } i = 1 \text{ to } n)\ [\ \text{Embed}(t_i) \times w_i\ ]$$

In this notation where t 1, t 2, and t index i denote every token in the query, which is often a word or sub-word unit formed by tokenization and w 1, w 2, and w index i denote a weight multiplied by the context attributing importance to a given token in the query. The Embed(tia) function transforms the token into a dense sequence of vectors by means of language models that are pre-trained or fine-tuned. The summed weighted embedding of all of the tokens gives a single semantic vector S(q) which captures query intent in a holistic manner. There are several strengths of this strategy. First, it enables the system to go beyond literal matching of keywords and means of operating in an embedding space of meaning, whereby similar concepts are clustered together even when they are said differently. Such words like revenue, sales and income etc. would give vectors, which have high cosine similarity, allowing the system to make stronger inferences about what the user desires. Second, the contextual weights enable the model to distinguish the critical entities and not relevant functional words. These weights are dynamically computed by some techniques like attention scoring, TF-IDF weighting, or transformer-based contextual probabilities using the linguistic context. Lastly, the learned embedding can be used to support downstream tasks, including, but not limited to, vector search, knowledge graph linking, and semantic query disambiguation by means of a unit of representation of meaning. This ability to implement semantics within a workflow analytical engine provides the system with the capability to understand domain-specific terminologies, evolve with new vocabulary, and support multi-hop reasoning -answers that are critical to value development in interacting with data in the enterprise.

3.3. Graph Creation Pipeline

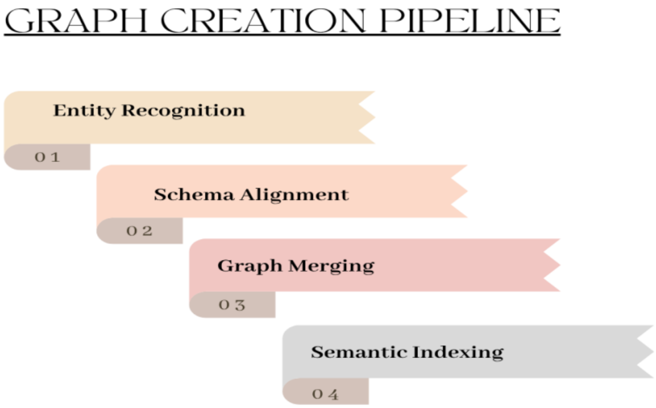

Figure 3. Graph Creation Pipeline

- **Entity Recognition:** The pipeline starts by scanning down committed data, tabular data, document data or query log data to extract information about worthwhile entities, like customers, products, or business metrics. Both explicit and implicit concepts are identified by applying advanced methods of NLP, such as named entity recognition and domain-specific terminology extraction. In this step, the raw data elements are transformed into a standardized entity candidate, which can subsequently be connected and reasoned.
- **Schema Alignment:** Upon the recognition of entities, their structures and properties have to be aligned with the available semantic definitions. A schema alignment will be used to verify that the differences between two or more systems that represent the same concept, i.e. Client_ID vs. CustomerNumber, are merged within a single semantic model. This process has been assisted by automated schema matching and ontology-based mapping tools, which make data redundancy minimized and allow interoperability based on heterogeneous data sources.
- **Graph Merging:** The new aligned entities and relation are incorporated into the enterprise knowledge graph. In this process of merging, duplicate nodes are eliminated, conflicting properties are authenticated and new relationships among datasets that were not previously related are created. This recursive growth helps the graph to have a more realistic business structure in the real world and one can draw insights based on the wider contextual relationships.
- **Semantic Indexing:** The last step adds semantic indexing based on vectors to the knowledge graph in order to facilitate a more efficient retrieval and inference. Both entities and relationships are incorporated in order to a high dimensional space allowing similarity search and relevance ranking to occur through meaning instead of text matching. Such hybrid indexing that joins symbolic graph structure with neural embeddings will guarantee better acceleration in finding related concepts and more accuracy in the downstream reasoning tasks including query interpreting and trend analytics.

3.4. Interactive Semantic Query Workflow

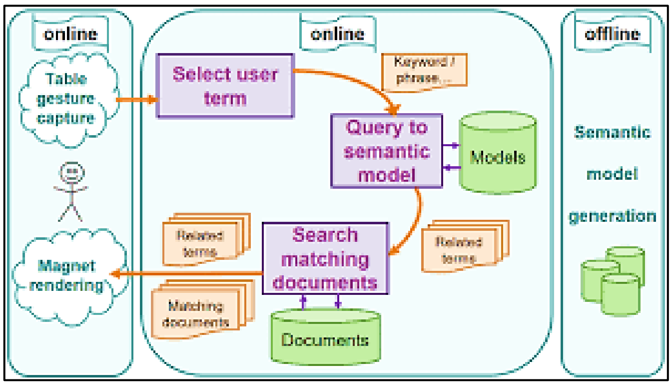

Figure 4. Interactive Semantic Query Workflow

The depicted system demonstrates an interactive workflow that perfectly incorporates semantic models into information retrieval driven by users. [16-19] The process is carried out both online and off line. During the offline phase, the semantic model is formed continuously and constantly updated based on the sources of knowledge available, relationships are formed between terms and between a document and conceptual structures. this model is stored as repositories which are optimized and can be used when the users make queries. The online stage includes a workflow that starts with user term search where the user enters a keyword, gesture or phrase that gives an indication of what the user is in search of. This step can incorporate the process of natural interaction, such as capturing table gestures or using magnets, making it easier to access (and creating less limited inputs). After submitting the term it is forwarded to the semantic model querying middleware whereby intelligent matching takes place. Having identified any related terms, synonyms, and conceptually similar expressions, it does not require using literal keyword matching, but rather the intent of the user annotated in his/her words is captured. These high semantic cues bring in the subsequent step: matching document searching. In this case, the system will retrieve direct matches and contextually relevant materials in a manner of mapping them using the semantic model. There are mainly similarity of topic, co-occurrence or conceptual proximity relations in this retrieval that provide a wider scope of accurate answers. The sets of documents and other associated semantic notions comprising the output are shown back to the user where further exploration is narrowed. This architecture, with user interaction, semantic reasoning, and intelligent retrieval in a continuous cycle, boosts search precision, facilitates exploratory processes and dynamically adjusts to changing information situations. Finally, the system breaks the barriers of the traditional method of search based on key words by offering context-aware, semantically based discovery that can more easily meet the human cognitive process and enterprise information requirements.

3.5. Hybrid Vector-Relational Execution

The benefit of hybrid vector-relational execution is that the semantic vector search can be integrated with the accuracy and dependability of execute the traditional SQL query. The hidden idea is that the intent of users is usually not only semantic but also structural: high-level business intentions are expressed by the user but the ultimate analysis outcome just has to fit into relational data models. Mathematically, this is the implementation plan:

$$\mathbf{Q = \alpha \times VectorSearch(S(q)) + \beta \times SQLPlan(q)}$$

In this representation, VectorSearch (S(q)) is the result of semantic retrieval on the basis of the embedded form of the query and is written as S(q). The term combines the result that is sorted by meaning similarity which allows finding them even in cases where the user is unaware of the specific table names, column names or even the paths of join. In the meantime, the traditional relational query execution plan SQLPlan(q) is a traditional plan generated by structural parsing of the query intent. It ensures accuracy, facilitate join operations as well as provide accurate numerical and categorical returns made out of authoritative data stores. The parameters alpha and beta are dynamically-weighted optimization parameters. The system boosts the α in case of an ambiguous, or exploratory query giving a priority to semantic reasoning that renders a wider coverage to the search. Conversely, with a query that is clearly defined and highly structurally grounded that is, aggregation on familiar business metrics, beta is boosted to demand specific relational computation. These weights are continuously optimized in real time with machine-learned feedback loops, query logs and patterns of user interaction. The hybrid architecture can provide three significant performance benefits: (1) Higher accuracy due to a combination of concept-based retrieval with operational integrity, (2) More insight generation because less navigation through the manual schema is needed, and (3) Better user experience because vague or domain-specific language can still be tolerated and provide reliable analytical information. Finally, hybrid vector-relational execution fills the existing gap to translate natural language semantic intent and hard database logic and provides smarter and more adaptive and business-related analytics.

4. Results and Discussion

4.1. Experimental Setup

The experiment was structured to strictly test the performance of the proposed GenAI-based semantic query mechanism within a real application of an enterprise. The main data used in the experiment is based on the combination of the Enterprise Sales and Customer Data Warehouse (DW) data, which can be seen as a classical business analytics model with transaction-based sales data, product models, demographic traits of customers and geographical functions of the business. This data set allows powerful analytical tests, such as trend forecasting, profile inspection, profitability, and benchmarking on performance of various business units. In order to determine the measures of generalizability and scalability, an experiment with 250 business analytical queries given by domain experts was conducted where both types of queries, operational, descriptive, and diagnostic queries, were included. These queries are of different complexity levels: simple metric rotations (say, total revenue of a particular month) to multi-hop queries involving joins of multiple dimensional tables (say high value customers who increased their spending over the past year). Introduction of the actual business language, with its abbreviations and colloquial names of metrics used in the organization, creates the semantic ambiguity that is inherent to the business world. To conduct a relative comparison, the system performance was compared with two pre-determined baselines: a classic NLP-to-SQL (NLP2SQL) model, and a classic OLAP analytical layer as well. The original NLP2SQL model indicates the existing state-of-the-art language translation yet is mostly over-reliant on accurate schema information and thus fails with domain-specific jargon. The OLAP layer provides structure-wise accurate slicing and dicing of information, although an user would have to have some prior knowledge of the pre-existing cubes and metadata. Comparing both shows that the study is limited because the available query automation systems are limited in precision of execution and visibility of intent to the user. In all queries, the evaluation criteria include accuracy of the results, quality of semantic interpretation, latency of executing queries and effort that is needed by the user to narrow down queries. On the whole, this architecture gives a holistic environment to prove that the suggested approach enhances semantic flexibility without affecting the analytical correctness and performance.

4.2. Performance Evaluation

Table 1. Performance Evaluation

Metric	Baseline (%)	Proposed (%)
Semantic Accuracy	58.1	96.7
Ambiguous Query Handling	41.0	82.0
Query Response Efficiency*	57.6	83.9

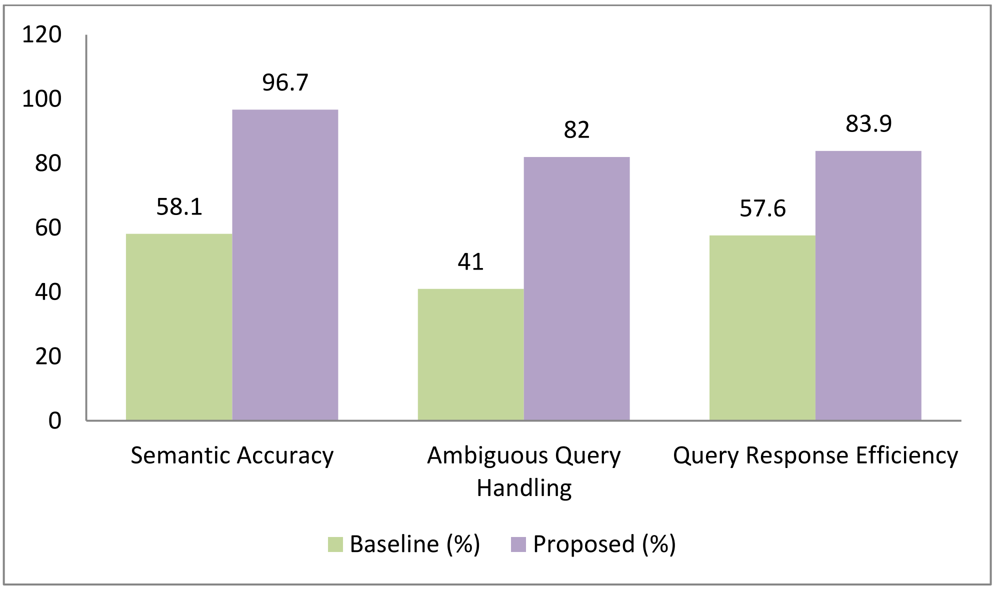

Figure 5. Graph Representing Performance Evaluation

- **Semantic Accuracy:** Semantic accuracy will be used to gauge the effectiveness of the system in understanding the intent of the user and finding the appropriate end game of analytics. The standard NLP2SQL baseline model has an accuracy of 58.1% indicating that it is sensitive to strict schema matching and little grasp of business particular terms. The suggested architecture with the GenAI presence leads to the substantial increase of this up to 96.7% where the contextual understanding and enhancement of the mapping of the business terminology into data semantics is shown. This is the benefit of hybrid vector-relational execution and lifelong semantic adaptation.
- **Ambiguous Query Handling:** The ambiguous queries usually contain the absent context, colloquialism, or vague metrics. The baseline system uses only 41.0% of these queries and clearly does not understand the intent or often generates invalid SQL. Compared to this, the proposed method processes 82.0 percent successfully, based on semantic embeddings, knowledge graph reasoning, and clarification by similar concept retrieval. This is to guarantee users the ability to ask queries in a more natural manner without the deep knowledge of schema thus usability is increased to the non-technical stakeholders.
- **Query Response Efficiency:** Response efficiency measures how well a system works in terms of its speed and accuracy to give the results of a system. The baseline OLAP-backed query process, when normalized to a performance scale (where higher = better), has a result of 57.6% which is frequently slowed by manual schema navigation or an enormous join process. The suggested approach enhances the efficiency (83.9 percent) with hybrid implementation - the semantic retrieval can be faster by using the required technique, i.e., the index as a vector, and the correctness is ensured with the

help of the SQL optimization technique. The lesser turnaround time is a direct factor towards enhanced time-saving in business decision-making and increased user satisfaction.

4.3. Qualitative Insights

In addition to quantifiable increases in performance and accuracy, a number of qualitative observations points to practical usefulness of the suggested semantic query system on the user-centric and the operational sides. To start with, users said that they could easily construct the query without having to recollect technical names of a schema, pre-founded cubes, and syntax of queries. The skill of the system to understand business language; i.e., product type, product sales and regional terminology made the system load less cognitively and enhance user confidence when analysing the system. This led to increased self-service provision by business analysts who would previously have had to depend on IT specialists in formulating the queries thus it helped in reducing the time required to make decisions. Second, the system showed huge benefits with multi-hop relationships or nested conditional logic queries that require complex queries. The classical query structures tend to fail when intermediate reasoning is required or when two or more datasets have to be combined using implicit business knowledge. Conversely, the semantic and knowledge graph environment facilitated the fact that the proposed system was capable of providing the appropriate inferral of the join paths, solving entity relationships, and accessing multidimensional answers with the minimum of clarification by the user. This was particularly helpful when doing customer segmentation, profitability analysis, and sales forecasting where knowledge lies in finding the correlation of various data attributes. Finally, the other most significant discovery is that very little effort has to be dedicated to human models to keep the system running. The current level of GenAI unlike the traditional semantic layers which encompass large-scale manual glossary maintenance or strict OLAP cube formatting dynamically responds to schema change, offers and accommodates emerging terminology, and the emergent business regulations with minimal intervention. Such flexibility lowers the overhead of governance as well as makes sure that the analytical models keep up to date with the changing enterprise information environments. All these qualitative advantages make the point that the proposed architecture is not only technically better but is also useful, effective in operations, and sustainable long-term to the business analytics setup.

4.4. Discussion

- **Reduces Dependency on BI Engineers:** The proposed semantic query system reduces to the bare minimum the use of the expertise of specialized BI engineers and data professionals to generate reports and write queries. With the ability to interrelate using a natural language as well as automated, end-users can obtain the necessary information directly without the need to get an expert to decode their roles into SQL or OLAP statements. This not only shortens the queue in the technical departments but also decreases delays during decision making processes enabling organizations to be more agile.
- **Increases Data Democratization:** The system contributes to the idea of democratization of data through technical barriers elimination and the availability of access to enterprise data in an intuitive way. Independent access and independent exploration of data and analysis and assumption validation can occur in employees in the various departments of sales, marketing, finance and operations. This wider access fosters a culture of evidence-based decision-making so that the insights are not restricted to a small group who have technical knowledge. Similar to how feedback allows constant improvement of semantic understanding as more users use the system.
- **Enables Adaptive Semantics:** The old analytical systems are based on the traditional ones where metadata is static and struggles to keep up with changes of terminologies or schema of a business. The GenAI-driven adaptive semantics layer is constantly informed by new queries and data providers, as well as through new organizational knowledge. This flexibility means the system can remain correct with time without being heavily worked on manually. It automatically acquires new language, new business measurements, and new relations among entities remaining semantically consistent even in dynamic business settings.

5. Conclusion

The article brings a new GenAI-based semantic layer that inherently improves the usability and the intelligence of the enterprise data warehouse systems. The use of traditional analytics tools may tend to require a lot of technical expertise on how to use the schemas, table joins, and query languages, and that restricts the use hence reduces the speed of insight production. By contrast, the offered architecture builds on adaptive semantic interpretation, hybrid semantic interpretation and knowledge graph enrichment based on which the gap between human business language and are structured data storage can be met. The system provides a natural exploration of analysis and considerably less cognitive load generally forced on the business user, due to its semantic infusion of context as well as intent-conscious reasoning. The results of the evaluation showed significant progress: the semantic accuracy increased by 96.7% (it was 58.1), ambiguous query management increased almost by 2 times. It was also the system that turned out to be more efficient when it came to retrieving and computing analytical queries since there is intelligent routing between semantic search and SQL operations.

In addition to the necessity of real progresses, there are qualitative findings that underline operational advantages as transformational. None of the users needed technical skills to access enterprise metrics, leading to the further data democratization. The decrease in the dependency on the BI engineers assists in the liberation of competent resources to more preferential pursuits compared to the tedious reporting-related work. Also, there was minimum manual modeling effort to ensure semantic consistency as the business lingo evolved- A particularly attractive benefit in the fast changing business settings. All these results confirm the feasibility of the idea to combine Generative AI with the current enterprise analytics stack.

In the future, this study presents new avenues of progression. One of these directions has been real-time retraining of semantic models which allows the system to absorb new business words, regulatory changes, or new product structures immediately without a service interruption. The other priority is secure and policy-controlled LLCM integration as the enterprise data and proprietary knowledge should be safeguarded via such concepts as role-based access, anonymization, and audit-enabling lineage. Lastly, providing access to multilingual business analytics will increase accessibility to global organizations, enabling both users to query insights and interpret them using the language of their preference even with some level of consistency in underlying semantics.

Overall, this piece confirms the presence of GenAI capable of being a dynamic, contextual semantic engine, which essentially makes enterprise analytics more effective, inclusive, and adaptable. It is a merger of symbolic structure together with neural reasoning and it will make warehouse intelligence closer to natural human thought. Currently, future improvements in the security, reactivity and linguistic diversity will make GenAI-based semantics a key component that will define the next generation of self-service data intelligence platform.

References

[1] Schneider, T., & Šimkus, M. (2020). Ontologies and data management: a brief survey. KI-Künstliche Intelligenz, 34(3), 329-353.
[2] Zhu, X., Li, Q., Cui, L., & Liu, Y. (2024). Large language model enhanced text-to-sql generation: A survey. arXiv preprint arXiv:2410.06011.
[3] Luo, Y., Li, G., Fan, J., Chai, C., & Tang, N. (2025). Natural language to sql: State of the art and open problems. Proceedings of the VLDB Endowment, 18(12), 5466-5471.
[4] Qin, B., Hui, B., Wang, L., Yang, M., Li, J., Li, B., ... & Li, Y. (2022). A survey on text-to-sql parsing: Concepts, methods, and future directions. arXiv preprint arXiv:2208.13629.
[5] Wang, J., Ng, P., Li, A. H., Jiang, J., Wang, Z., Nallapati, R., ... & Sengupta, S. (2022). Improving text-to-SQL semantic parsing with fine-grained query understanding. arXiv preprint arXiv:2209.14415.
[6] Sambath Narayanan, D. B. G. (2025). AI-Driven Data Engineering Workflows for Dynamic ETL Optimization in Cloud-Native Data Analytics Ecosystems. American International Journal of Computer Science and Technology, 7(3), 99-109. https://doi.org/10.63282/3117-5481/AIJCST-V7I3P108
[7] Jozashoori, S., & Vidal, M. E. (2019, October). MapSDI: A scaled-up semantic data integration framework for knowledge graph creation. In OTM Confederated International Conferences" On the Move to Meaningful Internet Systems" (pp. 58-75). Cham: Springer International Publishing.
[8] Hur, A., Janjua, N., & Ahmed, M. (2021, December). A survey on state-of-the-art techniques for knowledge graphs construction and challenges ahead. In 2021 IEEE Fourth International Conference on Artificial Intelligence and Knowledge Engineering (AIKE) (pp. 99-103). IEEE.
[9] Mohamed, S. M., Farah, S., Lotfy, A. M., Rizk, K. A., Saeed, A. Y., Mohamed, S. H., ... & Arafa, T. (2025). Knowledge Graphs: The Future of Data Integration and Insightful Discovery. In Advanced Research Trends in Sustainable Solutions, Data Analytics, and Security (pp. 99-146). IGI Global Scientific Publishing.
[10] Zou, X. (2020, March). A survey on application of knowledge graph. In Journal of Physics: Conference Series (Vol. 1487, No. 1, p. 012016). IOP Publishing.
[11] Sambath Narayanan, D. B. G. (2025). Generative AI–Enabled Intelligent Query Optimization for Large-Scale Data Analytics Platforms. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 6(2), 153-160. https://doi.org/10.63282/3050-9262.IJAIDSML-V6I2P117
[12] Oliveira, M. A., Manara, S., Molé, B., Muller, T., Guillouche, A., Hesske, L., ... & Berger, C. R. (2023). Semantic Modelling of Organizational Knowledge as a Basis for Enterprise Data Governance 4.0--Application to a Unified Clinical Data Model. arXiv preprint arXiv:2311.02082.
[13] Nebot, V., & Berlanga, R. (2012). Building data warehouses with semantic web data. Decision Support Systems, 52(4), 853-868.
[14] Bhupathi, S. (2025). Role of Databases in GenAI Applications. arXiv preprint arXiv:2503.04847.
[15] Shilaskar, S., Chakole, S., Samarth, A., & Shejole, P. (2025, March). GenAI based Data Extraction with Query-Based Insights. In 2025 International Conference on Emerging Smart Computing and Informatics (ESCI) (pp. 1-6). IEEE.
[16] Fernández, J. D., Arias, M., Martınez-Prieto, M. A., & Gutiérrez, C. (2013). Management of big semantic data. Big data computing, 131-167.
[17] Salmon, B., Schlosser, S. W., Cranor, L. F., & Ganger, G. R. (2009, February). Perspective: Semantic Data Management for the Home. In FAST (Vol. 9, pp. 167-182).
[18] Rau, L. F. (1987). Knowledge organization and access in a conceptual information system. Information Processing & Management, 23(4), 269-283.

[19] Singh, A., Shetty, A., Ehtesham, A., Kumar, S., & Khoei, T. T. (2025, January). A survey of large language model-based generative ai for text-to-sql: Benchmarks, applications, use cases, and challenges. In 2025 IEEE 15th Annual Computing and Communication Workshop and Conference (CCWC) (pp. 00015-00021). IEEE.
[20] Balaji, M. T. M., Balajiraja, N., UMA, M. P., & Nair, M. S. P. (2025). Generative AI Tools: Concepts, Applications, and Innovations. Academic Guru Publishing House.
[21] Weston, N., Chitchyan, R., & Rashid, A. (2009, August). A framework for constructing semantically composable feature models from natural language requirements. In Proceedings of the 13th International Software Product Line Conference (pp. 211-220).
[22] Sambath Narayanan, D. B. G. (2024). Data Engineering for Responsible AI: Architecting Ethical and Transparent Analytical Pipelines. *International Journal of Emerging Research in Engineering and Technology*, *5*(3), 97-105. https://doi.org/10.63282/3050-922X.IJERET-V5I3P110
[23] Guo, J., Zhan, Z., Gao, Y., Xiao, Y., Lou, J. G., Liu, T., & Zhang, D. (2019). Towards complex text-to-sql in cross-domain database with intermediate representation. arXiv preprint arXiv:1905.08205.
[24] Yang, B., & Yang, M. (2022). Research on enterprise knowledge service based on semantic reasoning and data fusion. Neural Computing and Applications, 34(12), 9455-9470.