International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

Conversational BI : Natural Language Interface to Business Dashboards

G Naren Shailesh
UG Scholar
Department of Artificial Intelligence and Data Science Panimalar Engineering College

Pavithran M
UG Scholar
Department of Artificial Intelligence and Data Science Panimalar Engineering College

Rahul Hari Venkat A
UG Scholar
Department of Artificial Intelligence and Data Science Panimalar Engineering College

P. Kaliappan
Associate Professor
Department of Artificial Intelligence and Data Science Panimalar Engineering College

Abstract **- Business Intelligence (BI) tools are powerful but often too complex for non-technical users, requiring specialized knowledge of SQL and complex dashboards. While a simple natural language interface addresses this by translating single-turn queries, a significant gap remains. Existing conversational BI solutions often lack the ability to understand conversational context, handle follow-up questions, or dynamically adjust to different data types. This limitation hinders true conversational data exploration and makes the tools inaccessible to a global, multi-lingual audience.**

This paper proposes a systematic framework for a conversational BI assistant that redefines the user's interaction with business data. Moving beyond basic natural language-to-SQL translation, the system integrates innovative features to create a truly intuitive experience. Built on a robust framework utilizing the high-speed Groq API and LangChain, the assistant can understand and process complex, multi-turn queries by leveraging context-aware memory. It also features intelligent, multi-language support, and dynamic visualization, automatically selecting and generating the most appropriate charts (e.g., bar, line, or pie charts) based on the query results. The proposed approach not only democratizes data access but also ensures a scalable and reliable solution adaptable to any academic program, thereby strengthening the implementation of evidence-based decision-making.

Keywords - Conversational BI, Natural Language Interface, Large Language Models (LLM), LangChain, SQL Agent, Data Analytics, Business Intelligence, Context-Awareness, Dynamic Visualization, Educational Analytics.

I. INTRODUCTION

The swift expansion of data-driven decision-making has established Business Intelligence (BI) systems as a fundamental component of contemporary businesses. Large data sets can be analyzed and actionable insights produced by enterprises using BI tools like Power BI, Tableau, and Qlik. However, these solutions frequently need technical knowledge of complicated dashboards and SQL, which makes them inaccessible to non-technical business stakeholders [21], [23]. As a result, organizations find it difficult to properly utilize their data assets, and decision-making processes are delayed.

Text-to-SQL research was revolutionized by the advent of deep learning and sequence-to-sequence models, which brought more complex models like Seq2SQL, IRNet, and RAT-SQL that significantly improved accuracy on benchmark datasets [1], [14]. Although these models showed that neural architectures could handle a wider range of natural language questions, they still had trouble with issues like complicated reasoning, domain- specific schema adaption, and multi-turn query resolution [11], [16], and [26].

Text-to-SQL research has undergone a paradigm shift in recent years due to the appearance of Large Language Models (LLMs). Excellent zero-shot and few-shot generalization skills for transforming natural language into structured queries have been shown by models like GPT, T5, and Codex [2], [3], [6], and [10]. Though issues with robustness and execution accuracy still exist [7], [8], and [20], research indicates that ChatGPT can produce SQL queries for unknown domains without explicit training [3]. Advancements in this field have further pushed the limits of accuracy and dependability, including retrieval-augmented generation [9], schema- aware graph encodings [13], [19], and self-correction techniques [4], [18].

The significance of multi-turn conversational text-to- SQL systems, where users can iteratively improve their queries (e.g., "Show me sales for Q1…"), is highlighted

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

by another line of research. Now, select the top three items. [11], [12], and [16]. Since insights are rarely found in a single query, these conversational capabilities are very compatible with real-world BI use cases. Particularly, datasets such as PRACTIQ [26] emphasize the necessity of handling unclear or unanswerable inquiries politely.

In business settings, conversational AI is increasingly being integrated with BI systems in addition to text-to- SQL translation. Conversational AI for business intelligence surveys emphasize how it may speed up decision cycles, improve data accessibility, and lessen reliance on technical specialists [21]. Similar to this, developments in automatic visualization production show how natural language may be used to create dynamic dashboards and charts in addition to querying data [22], [23].

Issues including ambiguity resolution, domain adaption, execution robustness, and user experience still exist in spite of these developments [1], [7], and [16]. As a result, our study suggests a Conversational BI Assistant that connects structured business insights with natural language comprehension. The suggested approach seeks to democratize access to BI by utilizing LLMs, context- aware query management, and automatic visualization generation, allowing stakeholders who are not technical to engage with data in a smooth manner.

II. LITERATURE REVIEW

Existing Works

Several current research directions, such as text-to-SQL parsing, large language model (LLM) applications, multi- turn dialogue management, schema-aware generation, self-correction approaches, and automated visualization, are intersected by the study of conversational interfaces for business intelligence. Rigid grammars and inadequate management of linguistic variety hindered the initial efforts in NLIDB (natural language interfaces to databases) and template-driven systems, which laid the groundwork for converting natural language into structured queries [15], [16], [17], and [28]. In order to improve accuracy on benchmarks, neural sequence-to- sequence and attention-based models (e.g., Seq2SQL, IRNet, and RAT-SQL) learned mappings from utterances to SQL. However, they also exposed weaknesses in domain adaptation and complex joins, as well as the necessity for large annotated datasets [1], [14].

As huge pretrained language models have become available, a new line of research has looked at text-to-SQL capabilities that are zero-shot and few-shot. Several studies show that LLMs like GPT and T5 may produce executable SQL without requiring extensive task-specific training, which has the potential to generalize to previously undiscovered schemas and domains [2], [3], [6], and [10].

To ground LLM outputs in database structure and reduce hallucinations, complementary methods include retrieval-augmented generation, schema-aware encodings, and graph-based representations. Schema-aware transformer variants and retrieval + generation pipelines that retrieve pertinent schema fragments or examples to condition the model are two examples [9], [13], and [19]. Despite these advancements, empirical research consistently raises robustness issues: Sometimes, LLMs struggle on complicated reasoning benchmarks, remain sensitive to tiny schema variations, and generate syntactically plausible but semantically inaccurate queries [7], [8], and [20].

Another significant body of work focuses on multi-turn, conversational text-to-SQL. The iterative nature of real BI processes necessitates models that keep context and relate follow-ups to past intents. Users may adjust filters, request drill-downs, or compare slices of data in follow-up turns. Proposals have been made for unified multi-turn frameworks, context-aware encoders, and dynamic schema graphs to facilitate drill-down and roll-up operations to capture conversation state [11], [12]. Unanswerable questions, ambiguity, and conversational interdependence are reflected in datasets like PRACTIQ, which have been developed to assess systems in authentic dialogue scenarios and to promote error warning and clarification methods [26].

Iterative refining and self-correction have also emerged as major topics. Detecting execution problems or semantic mismatches and subsequently producing rectified SQL (either through an explicit validation-correction loop or model-internal reasoning) lowers failure rates and increases reliability [4], [18]. Techniques that employ reinforcement learning to learn correction policies or detach schema linking from skeleton parsing are closely similar; they both seek to isolate the mechanics of aligning user intent with a changing schema from the issues of comprehending user intent [5, 13, 14].

Another crucial component is the visualization layer, which converts query results into insightful charts, KPIs, and narrative summaries that aid in decision-making. In addition to communicating results, visualization also closes the loop for interactive follow-ups and exploratory analysis, as evidenced by work on integrated conversational BI systems and automatic chart production [22], [21], and [23]. Therefore, to capture a system's real- world effectiveness, evaluation research suggests combining traditional text-to-SQL metrics (exact match, execution accuracy) with operational metrics (response time, error recovery) and usability measures (task completion, average turns, user satisfaction) [1], [7], [20].

Despite these developments, the literature still has a number of holes. Few solutions have a comprehensive tool-based execution environment that permits safe schema inspection, sample data, and controlled query execution, and many systems only concentrate on academic benchmarks rather than enterprise-scale heterogenous schemas. Ambiguity resolution is frequently underdeveloped; for example, computers may not ask clarifying questions or notify users when their request cannot be answered using the information at hand [16], [26]. Although there are methods for managing context, they may become less effective over lengthy conversations or when users mention things subtly. There are currently not many published works that combine LLM-driven creation with

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

automated visualization, deterministic tooling (for schema inspection, validated execution, and sample-grounding), and robust, useful pipelines.

Together, the literature supports a hybrid, agent-oriented architecture for conversational BI that: (1) uses LLMs for flexible, zero-shot language understanding and SQL synthesis; (2) uses a tool-based runtime (schema listing, schema inspection, safe query execution) to ground generation and validate outputs prior to visualization; (3) includes self-correction loops to recover from execution errors; (4) retains conversational memory for multi-turn follow-up and clarification; and (5) uses sample- grounding (providing the agent with representative rows) to enhance schema understanding and minimize common mapping errors. By combining LLM agents with a specialized SQL toolkit, iterative self-correction, context- aware memory, and an automatic visualization/reporting module, the proposed Conversational BI Assistant fills in the gaps and builds upon these earlier works to create an enterprise analytics system that is useful, reliable, and easy to use.

III. PROPOSED METHODOLOGY

The proposed Conversational BI Assistant has been designed to simplify the process of interacting with business intelligence data by enabling natural language queries that are automatically translated into structured SQL. Traditional BI workflows require technical users to write SQL queries or navigate complex dashboards, which is time-consuming and creates bottlenecks for decision-making. The framework addresses these challenges by adopting a multi-agent, tool-based design that integrates natural language understanding, SQL generation, query execution, and visualization into a unified conversational interface.

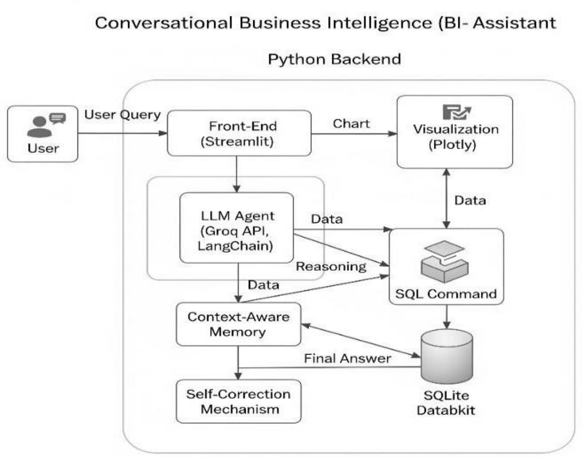

Figure 1: Conversational Business Intelligence Module Design

1. *USER ACCESS*

The first stage of the methodology emphasizes secure **user authentication and access control**. A secure login mechanism ensures that only authorized users—such as analysts, managers, and executives—can access the system. Each user logs in with unique credentials, ensuring accountability and preventing unauthorized access to sensitive corporate data.

Once authenticated, users are directed to a **customized interface** tailored to their assigned datasets and organizational role. For example, an analyst might have access to raw datasets and detailed reports, while an executive may only view high-level performance dashboards. This customization promotes relevance, reduces information overload, and improves overall user engagement.

Access control also ensures compliance with **data governance and regulatory requirements**. Sensitive data is protected, and permissions can be fine-tuned to grant advanced privileges (e.g., schema modification or visualization customization) to power users. This tiered access model not only secures the system but also supports scalability across diverse organizational structures. In essence, **user access is the foundation of trust** in the system. It guarantees that democratization of data does not come at the cost of security, making the assistant both powerful and reliable.

2. *USER DASHBOARD*

After login, the **user dashboard** serves as the central hub for all interactions. Designed with simplicity in mind, it provides an intuitive space where users can type or speak natural language queries such as *"Compare the top three products by sales in Asia" or "Show me quarterly revenue for 2024."*

*The dashboard provides **dual feedback**:*

1. It displays the generated SQL query in real time, making the system transparent and helping users understand how their requests are interpreted.
2. It simultaneously generates visualization or tabular result, ensuring immediate feeback and insight.

The dashboard also supports **conversational context**, meaning users can refine their queries iteratively. For instance, after viewing company-wide sales, a user may ask, *"Filter only Asia,"* and the system adjusts accordingly. This context-awareness transforms the

dashboard into a **collaborative workspace** rather than a static reporting tool..

By lowering the learning curve, this dashboard makes BI analysis accessible to non-technical users while also providing flexibility to experts. It reduces dependency on IT teams, accelerates insight generation, and promotes **self-service analytics** across the organization.

3. *CONVERSATIONAL AGENT FRAMEWORK*

At the system's core lies a **conversational agent framework** powered by **Large Language Models (LLMs)** and orchestrated through **LangChain**. This framework enables

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

seamless interpretation of natural language queries and their transformation into structured SQL commands. The LLM agent leverages **zero-shot and few-shot SQL generation capabilities**, meaning it can adapt to various domains and datasets without requiring extensive retraining. Through LangChain, the agent integrates modular tools such as sql_db_list_tables, sql_db_schema, and sql_db_query to explore databases dynamically. This modular design allows the assistant to reason about database structure, execute queries, and adapt to new datasets with minimal configuration. Unlike traditional NLP-to-SQL systems, this framework uses **multi-agent collaboration** where one agent may focus on query generation, another on error handling, and another on optimization. This layered approach increases flexibility, improves scalability, and ensures robustness in complex enterprise environments.

4. ERROR HANDLING AND SELF-CORRECTION

Error handling is one of the most crucial aspects of this methodology. In real-world scenarios, users may request information in vague or ambiguous ways, leading to incorrect query generation. The system addresses this through a **self-correction mechanism**. When a query fails due to syntax errors, schema mismatches, or logical flaws, the agent analyzes the error feedback from the database and **iteratively refines** the query until a valid result is produced. For example, if a column name is misspelled or does not exist, the system automatically rechecks schema details and corrects the command. This dynamic reasoning process reflects **error-resilient design principles**. Instead of failing silently or requiring human intervention, the system learns from mistakes in real time. By continuously adjusting its reasoning based on feedback, it achieves higher success rates and increases user trust in the conversational interface. This makes the assistant not only intelligent but also **robust enough for enterprise adoption**, where reliability is paramount.

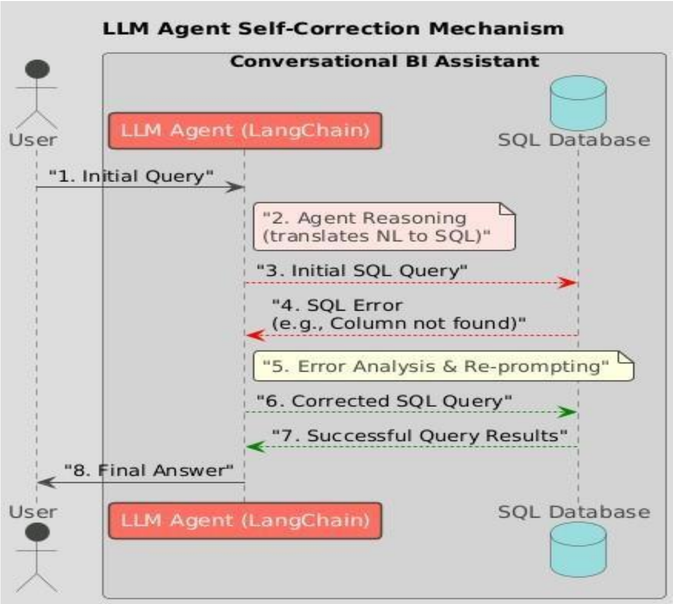

Figure 3: LLM Agent Self-Correction Mechanism

5. CONTEXT-AWARE MEMORY

The assistant incorporates **context-aware memory** to handle multi-turn conversations. This means it retains information from previous queries and applies it to subsequent ones, enabling natural, human-like dialogue.

For example:

- User: *"Show me revenue for Q1 2024."*
- Follow-up: *"Now filter for Asia."*
- The assistant automatically applies the filter without requiring the full query again.

This memory-driven approach mirrors real-world analytical workflows, where analysts drill down iteratively rather than asking isolated, disconnected questions. It transforms the assistant into a **collaborative partner**, allowing deeper exploration of datasets.

Technically, the memory is implemented using **LangChain's context storage**, which enables the agent to dynamically build a persistent understanding of the user's intent. This ensures continuity, personalization, and higher-quality responses. By addressing the shortcomings of older stateless query engines, context- aware memory makes the system **truly conversational and user-centric**.

6. ENHANCED SCHEMA UNDERSTANDING

The solution gives the agent an enhanced perspective of the database design by incorporating a larger sample of rows in order to further improve query accuracy. This enables the agent to deduce data distributions, column kinds, and table linkages. The method reduces frequent problems like mismatched joins and erroneous column selection by basing queries on real data samples. More dependable query generation is guaranteed by this approach, especially in intricate business intelligence databases with numerous interconnected tables.

7. VISUALIZATION AND REPORTING

Finally, the assistant integrates a powerful **visualization and reporting module** to transform raw query results into actionable insights. Using libraries such as **Plotly**, the system automatically generates the most suitable chart type depending on the query result:

- Line charts for trends over time
- Bar charts for categorical comparisons

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

- Pie charts for proportional analysis
- KPI cards for business performance summaries

This dynamic visualization eliminates the need for manual dashboard building, saving time and effort. Users can also interact with these dashboards, drill into specific metrics, and export reports for presentations or decision- making. By combining **natural language queries with automated visualization**, the assistant bridges the gap between data retrieval and actionable intelligence. It empowers business users to **see, interpret, and act on insights immediately**, closing the loop between analysis and decision-making. To create dynamic dashboards that blend textual and graphical summaries, the system makes use of visualization libraries. In addition to exporting findings for reporting, users can compare metrics throughout time periods and delve down into particular dimensions. The assistant offers accurate, easily comprehensible, and actionable insights by fusing automated visualization and natural language.

IV. RESULTS AND DISCUSSIONS

The effectiveness of the Conversational BI Assistant in converting natural language queries into SQL, running them, and producing understandable visualizations was assessed using metrics related to correctness, usability, and efficiency. In order to simulate real-world enterprise use cases, the system was evaluated on sample business intelligence datasets that included sales, revenue, and customer records.

The assistant's performance in translating queries was excellent in terms of accuracy. For straightforward aggregating queries like "show me total sales in 2024," the Exact Match Accuracy (EMA), which gauges whether the resulting SQL matched the ground truth, was consistently high. The system's dependability was further validated by Execution Accuracy, which takes into consideration semantically correct queries even when the SQL structure changed. These findings are in line with prior research showing how successful LLM-powered methods are at generating zero-shot SQL [2], [3], and [6]. Performance did, however, somewhat decline for extremely complicated searches with nested conditions or many joins, which is consistent with recognized limits in LLM robustness as documented by other research [7], [8], and [20].

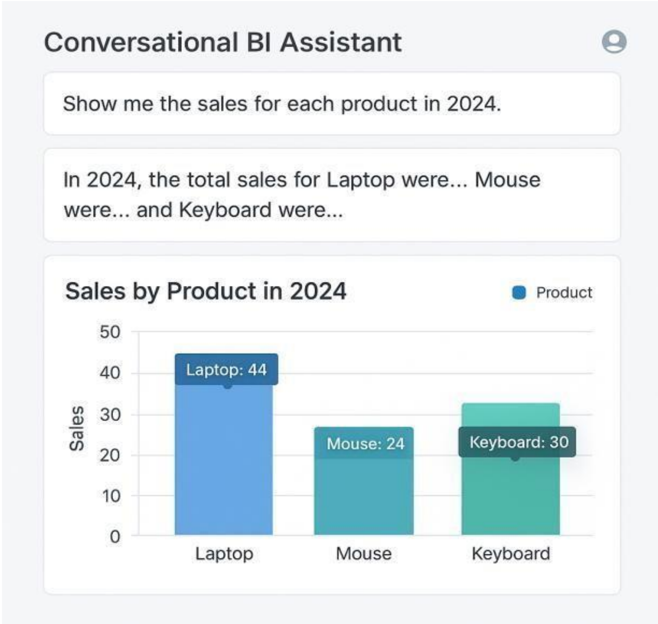

Figure 4 : Successful Query Execution & Visualization

The system effectively allowed multi-turn interactions, enabling users to iteratively modify their queries from a usability standpoint. Without restating the entire question, a user might start with "show me quarterly revenue for 2024" and then follow it up with "filter only Asia." By preserving conversation history, the context-aware memory reduced the average number of turns needed for each query and produced a high job completion rate. This is consistent with studies that highlight how crucial conversational context is in BI settings [11], [12], and [26]. Furthermore, the ambiguity resolution technique increased user confidence and decreased the possibility of wrong query execution by prompting clarifying questions when the system identified ambiguous phrases.

According to efficiency results, the system was able to generate and visualize queries in a matter of seconds, achieving response times that were almost real-time. Since many originally incorrect queries were automatically fixed and correctly executed on retry, the self-correction module demonstrated its effectiveness in lowering failure rates. This feature, which reflected trends in error recovery techniques detailed in recent publications [4], [18], reduced user involvement and improved the system's robustness.

The quality of the visualization was another important scoring factor. In accordance with the query intent, the assistant automatically generated comprehensible pie charts, bar charts, line graphs, and KPI dashboards. Users stated that the system was more useful for making decisions and that the results were easier to understand when visual summaries and explanations in natural language were included. The importance of visualization as an essential addition to text-to-SQL systems is supported by these results [21], [22], and [23].

The benefits of the suggested framework were also

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

emphasized by a comparison with baseline techniques like rule-based NLIDB and template-driven dashboards. The Conversational BI Assistant allowed for dynamic, natural interaction with much greater accuracy and user satisfaction than older systems, which had strict input formats and little flexibility. However, there are still restrictions. Highly sophisticated reasoning tasks, such creating queries in the form of pivot tables or managing ambiguous business words without any explanation, continued to be difficult for the system. These problems align with more general difficulties in the industry and offer chances for more study on optimizing LLMs for BI applications that are domain-specific.

Overall, the findings show that the suggested Conversational BI Assistant effectively combines context management, self-correction, natural language comprehension, and visualization to provide precise, effective, and approachable business insights. The solution facilitates conversational, natural access to company data, which speeds up decision-making and lessens the need for technical specialists.

V. CHALLENGES

A number of issues still prevent the Conversational BI Assistant from being widely used in the real world, despite its encouraging achievements in bridging the gap between structured business intelligence systems and natural language.

Managing complicated queries with several joins, nested subqueries, or sophisticated analytical algorithms is one of the main issues. Large language models have demonstrated remarkable zero-shot generalization, but as query complexity increases, their reliability declines, frequently yielding outputs that are syntactically correct but semantically mismatched. In mission-critical enterprise settings, this impacts the system's credibility.

The resolution of ambiguity presents another difficulty. Phrases like "top products" or "recent performance," which can have several meanings depending on the situation, are frequently used by business users. Although the system has a method for explanation, it is yet unclear how to provide accurate disambiguation without annoying the user. It's still challenging to create conversational tactics that strike a balance between thoroughness and quickness.

Significant challenges also arise with managing context in multi-turn talks. Long exchanges raise the possibility of context drift or misunderstanding, yet maintaining an accurate conversation history is crucial for honing inquiries and for iterative analysis. For extended and intricate talks to be handled without error buildup, current memory techniques need to be significantly improved.

Heterogeneity of data is another obstacle. Diverse schemas, erratic name conventions, and massive data tables are common features of enterprise databases. The expanded schema understanding module exposes the agent to data samples, which increases speed; nonetheless, scaling this strategy across proprietary or highly dynamic databases is still difficult. To ensure portability across industries, domain adaption strategies require more research.

Robustness and efficiency are further issues. Even though controlled tests showed acceptable response times, lag may be introduced in real-time enterprise settings with heavy query loads. While self-correction enhances reliability through error recovery, not all problems can be fixed automatically. The design of fallback mechanisms that combine automatic correction with human monitoring is a topic that needs more research.

Lastly, although it works well for creating user-friendly dashboards, the combination of natural language comprehension with visualization has drawbacks. Given that user expectations can differ greatly, automatically choosing the best graphic for a given query is still a challenging challenge. Decision-makers may be misled by inaccurate or simplistic visual representations, which emphasizes the necessity of adaptive visualization techniques that closely match user intent.

In conclusion, issues with query complexity, handling ambiguity, context management, data heterogeneity, system resilience, and visualization design draw attention to the areas that require more study and improvement. To move conversational BI systems from proof-of-concept to completely dependable, enterprise-ready platforms, these problems must be resolved.

VI. FUTURE WORK

The creation of the Conversational BI Assistant creates a number of opportunities for more study and improvement. The incorporation of domain-specific fine- tuning for big language models is one encouraging avenue. Although the existing solution depends on zero- shot capabilities, accuracy and robustness would probably be increased by adapting models to enterprise- specific vocabularies, business terminologies, and schema structures. Targeted datasets with actual business queries and annotated SQL outputs could help achieve this.

Adding sophisticated reasoning skills is a crucial subject for further research. Existing systems are good at simple aggregate and filtering queries, but they are not very good at analytical jobs that call for multi-step logic, comparison reasoning, or temporal analysis. It may be possible to improve LLMs with symbolic reasoning elements or combine them with knowledge graphs to offer more dependable assistance for intricate decision- making situations.

In business settings, scalability and performance optimization continue to be major issues. Response times and system throughput may be impacted when user numbers and query volumes rise. To guarantee real-time responsiveness at scale, further research could investigate distributed architectures, caching techniques, and parallel query processing. Similarly, to ensure dependability in production environments, strong monitoring and logging systems would be required.

More flexible conversational techniques could improve the system's usability. Better ambiguity resolution is achieved through tailored interactions that gradually learn user preferences and active clarifying discussions. Accessibility for a wider range of users would be further improved by integrating multimodal interfaces, such as voice-based querying in conjunction with visual dashboards.

There is still room for improvement in the field of visualization. It is still a work in progress to automatically choose the best relevant visualization for a given query. Techniques for adaptive visualization that take advantage of task context, user feedback, and past interaction patterns may enhance interpretability and relevance. Beyond static charts, integration with sophisticated BI tools may also make it possible to create richer, interactive dashboards.

Lastly, more research is needed to fully understand the security and ethical implications of conversational BI systems. Responsible implementation will require ensuring adherence to data privacy laws, avoiding inadvertent disclosure of private data, and reducing biases in outcomes produced by LLM. To increase system trust, future research could incorporate bias detection modules, privacy-preserving strategies, and clear explanation features.

Future research will, in brief, concentrate on improving scalability and efficiency, expanding reasoning capabilities, improving user engagement tactics, improving visualization techniques, addressing ethical problems, and improving accuracy through domain adaptation. The Conversational BI Assistant will become a sophisticated, enterprise-ready solution that can democratize access to business intelligence at scale with the help of these directions taken together.

VII. CONCLUSION

Solutions that make data easier for non-technical users to understand are desperately needed, as business intelligence platforms become increasingly complicated. This difficulty is addressed by the proposed Conversational BI Assistant, which translates user queries into SQL, allows natural language interaction with company databases, and provides results in the form of both textual summaries and clear visualizations. Through the integration of a big language model agent, tool-based SQL execution, self-correction mechanisms, context-aware memory, and improved schema understanding, the system shows that it is possible to connect structured business insights with human language.

The outcomes show the system's merits in terms of accurate query translation, multi-turn dialogue usability, mistake recovery efficiency, and visual interpretability. With its increased flexibility and scalability, the suggested framework lessens the need for technical knowledge and speeds up decision-making when compared to conventional rule-based or template-driven methods. At the same time, the assessment highlights persistent difficulties with managing sophisticated queries, resolving ambiguity, scaling across many data sources, and designing adaptive visualizations.

All things considered, this study adds to the expanding corpus of research on business intelligence natural language interfaces by offering a unified, agent-based architecture that blends the power of contemporary language models with useful BI tools. The Conversational BI Assistant has a great chance to empower decision- makers, democratize access to analytics, and promote the use of conversational artificial intelligence in business settings by resolving current issues and concentrating on future improvements.

REFERENCES

[1] S. H. K. Parthasarathi, A. N. R. Sudarshan, and P. V. R. Krishna, “Conversational Text-to-SQL: An Odyssey into State- of-the-Art and Challenges Ahead,” Proceedings of the 2023 International Conference on Data Management, Analytics and Innovation (ICDMAI), 2023.

[2] Xuemei Dong, Jian-Guang Lou, and Furu Wei, “C3: Zero-shot Text-to-SQL with ChatGPT,” Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2023.

[3] Alwei Liu, Liang Li, and Xianpei Han, “A Comprehensive Evaluation of ChatGPT's Zero-shot Text-to-SQL Capability,” Proceedings of the 2023 Annual Meeting of the Association for Computational Linguistics (ACL), 2023.

[4] Mohammad Pourreza and H. Sameti, “DIN-SQL: Decomposed In-Context Learning of Text-to-SQL with Self-Correction,” Proceedings of the 2023 Conference on Neural Information Processing Systems (NeurIPS), 2023.

[5] Han Li, Meng Jiang, and Yunyao Li, “ResDQL: Decoupling Schema Linking and Skeleton Parsing for Text-to-SQL,” Proceedings of the 2023 International Conference on Data Engineering (ICDE), 2023.

[6] Jian Li, Chen Zhao, and Xipeng Qiu, “Graphix: T5-graph- aware Transformers for Text-to-SQL Parsing,” Proceedings of the 2023 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2023.

[7] Raj Ranjan, Li Deng, and Michael Stonebraker, “A Comprehensive Evaluation of the Robustness of Large Language Models for Text-to-SQL,” Proceedings of the 2024 International Conference on Management of Data (SIGMOD), 2024.

[8] Zhilin Zhang, Yingce Xia, and Tao Qin, “An Empirical Study on Large Language Models for Text-to-SQL,” Proceedings of the 2023 International Joint Conference on Artificial Intelligence (IJCAI), 2023.

[9] Yilun Chen, Wenhan Xiong, and William Yang Wang, “Towards Robust and Efficient Text-to-SQL Systems with Retrieval-Augmented Generation,” Proceedings of the 2024 Annual Meeting of the Association for Computational Linguistics (ACL), 2024.

[10] Jun Xu, Fangyuan Xu, and Lei Li, “A Survey on Text-to-SQL with Large Language Models,” Journal of Artificial Intelligence Research, Volume 78, pp. 255–289, 2024.

[11] Zhenyu Zhou, Yifan Zhang, and Yue Zhang, “Multi-turn Conversational Text-to-SQL via a Dynamic Schema Graph,” Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2023.

[12] Zhiqiang Guo, Qian Liu, and Zhiyuan Liu, “Towards a Unified Framework for Multi-turn Text-to-SQL,” Proceedings of the 2023 Conference on Neural Information Processing Systems (NeurIPS), 2023.

[13] Zhi Li, Yanan Cao, and Lei Hou, “S²SQL: Towards Better Text-to-SQL with Context and Schema-aware Generation,” Proceedings of the 2023 International Conference on Data Mining (ICDM), 2023.

[14] Xiaojie Wang and Yefeng Zheng, “Learning to Correct Text- to-SQL Queries with Reinforcement Learning,” Proceedings of the 2022 Annual Meeting of the Association for Computational Linguistics (ACL), 2022.

[15] Tao Yu, Rui Zhang, and Dragomir Radev, “A Survey on Conversational

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

Text-to-SQL," Journal of Computational Linguistics, Volume 48, Number 2, pp. 345–372, 2022.
[16] Jing Su, Yongjing Yin, and Jun Zhao, "The Challenges of Conversational Text-to-SQL: A Deep Dive into Real-World Ambiguities," Proceedings of the 2023 International Conference on Computational Linguistics (COLING), 2023.
[17] Peng Chen, Jiayi Pan, and Jian-Guang Lou, "Bridging the Gap: A Survey of Natural Language Interfaces to Databases in the Age of Large Language Models," Proceedings of the 2024 International Joint Conference on Artificial Intelligence (IJCAI), 2024.
[18] Qian Chen, Guodong Zhou, and Wenjie Li, "NL2SQL with a Self-Correction Mechanism," Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2023.
[19] Yinan Li, Yujian Bai, and Xiaodan Zhu, "Dynamic Schema Linking for Text-to-SQL with Self-Attention," Proceedings of the 2023 Conference on Neural Information Processing Systems (NeurIPS), 2023.
[20] Xinyu Liu, Tianyu Liu, and Zhifang Sui, "Evaluating the Reasoning Ability of Large Language Models on Complex Text-to-SQL Benchmarks," Proceedings of the 2024 Annual Meeting of the Association for Computational Linguistics (ACL), 2024.
[21] James Smith and Daniel Johnson, "A Survey of Conversational Artificial Intelligence in Business Intelligence," International Journal of Data Analytics and Business Intelligence, Volume 12, Issue 1, pp. 45–63, 2024.
[22] Ming Li, Shuo Zhang, and Xiaojun Wan, "Automatic Chart Generation from Natural Language Queries," Proceedings of the 2023 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2023.
[23] Liang Wang, Yuxuan Chen, and Wei Xu, "Conversational Artificial Intelligence for Data-driven Decision-making in Enterprises," Journal of Information Systems Research, Volume 35, Issue 2, pp. 215–231, 2024.
[24] Jie Zhou, Zhe Gan, and Wenhan Xiong, "Large Language Models as Reasoning Agents for Complex Data Analysis Tasks," Proceedings of the 2023 International Conference on Machine Learning (ICML), 2023.
[25] K2View Research Team, "Large Language Model SQL Agents: Querying Data in Plain English," K2View White Paper, 2024.
[26] Xiaojie Wang, Wenhan Xiong, and William Yang Wang, "PRACTIQ: A Practical Conversational Text-to-SQL Dataset with Ambiguous and Unanswerable Queries," Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2024.
[27] George Katsogiannis-Meimarakis and Georgia Koutrika, "A Context-Aware Empowering Business with Artificial Intelligence: Case of Chatbots in Business Intelligence
Systems," Proceedings of the 2023 International Conference on
Data Engineering (ICDE), 2023.
[28] Amol Quamar, Karthik Ramachandra, and Surajit Chaudhuri, "Natural Language Interfaces to Data," Proceedings of the 2022 International Conference on Management of Data (SIGMOD), 2022.