International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V3I3P109
Eureka Vision Publication | Volume 3, Issue 3, 81-89, 2022

Original Article

Natural Language Querying in Oracle Fusion Analytics: A Step toward Conversational BI

Partha Sarathi Reddy Pedda Muntala
Independent Researcher, USA.

Abstract - Natural Language Processing (NLP) has emerged as a transformative technology in Business Intelligence (BI), enabling business users to query complex data systems without requiring specialised technical knowledge. Oracle Fusion Analytics is a modern cloud-based analytics suite that integrates seamlessly with Oracle Cloud Applications, offering deep insights across domains like finance, HR, supply chain, and customer experience. By integrating natural language querying (NLQ) capabilities into Oracle Fusion Analytics, enterprises can significantly enhance user accessibility, data democratization, and decision-making speed. This paper provides a comprehensive study of the integration of NLP interfaces in Oracle Fusion Analytics with a focus on Conversational BI. We present an in-depth analysis of how natural language interfaces, including text-based and voice-based assistants, empower users to generate self-service reports and obtain real-time answers to business questions. Furthermore, we examine the technological foundations of NLP in analytics, explore existing literature, propose a methodology for implementing NLQ in Oracle Fusion Analytics, and discuss experimental results. Key challenges such as language ambiguity, user intent detection, and performance optimization are addressed. This paper aims to contribute to the growing body of research by presenting a structured approach, supported with figures, tables, flowcharts, and real-world application scenarios, highlighting the future potential of conversational BI.

Keywords - Oracle Fusion Analytics, Natural Language Processing, Conversational BI, Voice Assistants, Business Intelligence.

1. Introduction

1.1. The Evolution of Business Intelligence

Business Intelligence (BI) has undergone significant changes in the last few decades, evolving from passively driven reporting systems to IT-driven, dynamic user-oriented platforms that enable decision-making at all levels within an organisation. The early development of BI was mostly centred around structured reporting, where the extraction and analysis of data were controlled within a specific group of IT specialists. Still, they had to use complicated SQL queries and data warehouses. [1-4] Although these systems were very powerful, they lacked flexibility and could not be accessed by non-technical users. With the further development of technology, the emergence of interactive dashboards and business intelligence self-service tools has allowed business users to navigate interesting data via visual front ends, such as Tableau and Power BI, and Oracle analytics. This has changed the landscape, as users no longer need advanced technical knowledge to examine data. Artificial intelligence (AI) and machine learning (ML) have become the next significant milestones in the history of BI, as they opened the door to the introduction of predictive analytics, anomaly recognition, and auto-insights. The involvement of natural language processing (NLP) is a relatively recent development that has started redefining the way users interact with data, as it enables users to ask questions in everyday language and receive relevant answers. This transformation is an outgrowth of a larger movement toward democratising access to data, reducing the technical division of labour, and democratising data analysis by making it easier, more approachable, and more generally understandable. The overall objective of contemporary BI is to translate data into actionable insights in a format that is instant, precise, and familiar to the user, in the context of how they talk, make decisions, and communicate.

1.2. Need for Conversational Interfaces in BI

- **Bridging the Gap between Users and Data:** With organizations growing more data-powered, the capacity for business users to access and decipher data in an efficient manner has become important. The traditional BI tools, though highly effective, usually require users to know the dashboards, browse through filters, or query. This presents an obstacle to other non-technical stakeholders who may not be competent in query languages and data analysis tools. A conversational interface can resolve this dilemma because it enables non-technical users to communicate or deliver voice or typed speech to interact with BI systems, thereby reducing the barrier of entry and making data more accessible to a wider range of users.
- **Enhancing Self-Service and Agility:** Businesses today require speed in decision-making, and one might have to wait until IT teams execute their reports or dashboards. Conversational BI interfaces enable users to pose questions in real-time, such as "What were sales last quarter in the North region?" and receive an immediate response. This is an

improvement in self-service analytics that enables quicker and more informed decisions, without requiring technical reintermediation. It reduces the data team's workload, allowing them to focus on more strategic projects.

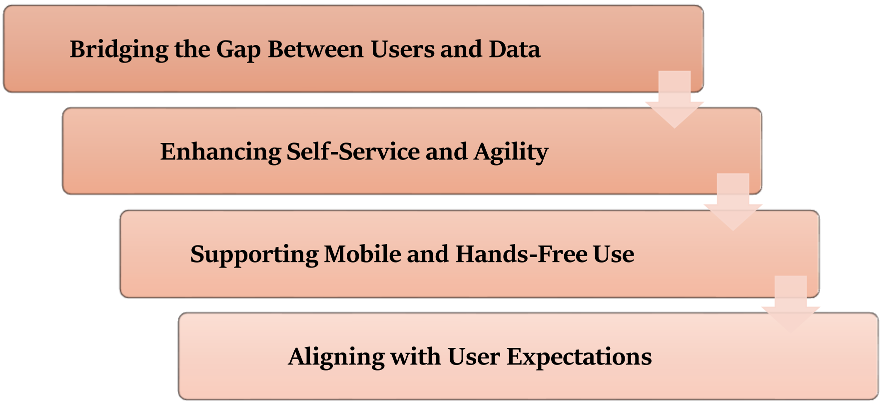

Figure 1. Need for Conversational Interfaces in BI

- **Supporting Mobile and Hands-Free Use:** As the remote working pattern and mobile business increase, the need for BI tools with on-the-field applications is also on the rise. These voice-enabled conversational UI interfaces can be part of a mobile application or a smart device that enables executives and field workers to access data by multitasking without typing or getting lost in a complex UI. This is hands-free, which enhances access and productivity in quick-paced or mobile-intensive positions.
- **Aligning with User Expectations:** With the introduction of technologies such as Alexa, Siri, or Google Assistant, and the simplicity of the communication flow they have created, users find it harder to work with tools that do not offer them such ease of communication in the workplace. Those expectations are consistent with the nature of conversational BI that makes data interaction feel as simple as asking a question to a colleague. Such a humanistic approach to design enhances the user experience, increases adoption rates, and ultimately leads to a more data-driven culture within an organisation.

1.3. Oracle Fusion Analytics Overview

Oracle Fusion Analytics is a next-generation business intelligence and analytics solution, engineered to integrate with Oracle Cloud Applications, including ERP, HCM, SCM, and CX. It is built to provide actionable and comprehensive insights, pre-hashed with prebuilt dashboards, KPIs, and over 1,300 predefined metrics by business function. These breakthrough analytics save on set-up time and enable organizations to realise the instant value of their data. The platform has closed integration with the transaction systems of Oracle and guarantees real-time synchronization and consistency of data between reporting and operational environments. One of the advantages of Oracle Fusion Analytics is that they are extensible and built with machine learning. End users have the freedom to add personal data models, algorithms, and representations to prebuilt analytics, tailoring them to meet the enterprise's requirements. The integrated ML capabilities enable predictive analytics, trend analysis, and anomaly detection, allowing commercial enterprises to transition to prescriptive decision-making.

These attributes qualify the platform for widespread adoption in various industries, including finance, healthcare, retail, and manufacturing. Oracle Fusion Analytics is also very popular due to its modularity, scalability, and enhanced data governance, which makes it suitable for interconnection with natural language processing (NLP). The application of NLP to the platform enables business users to access information in complex datasets using conversational interfaces, such as text and voice, without needing to know how to write queries or navigate dashboards. This significantly boosts interactivity, as users and analytics become more natural, and we no longer depend on IT or data experts. Whether the user wishes to pose a question, such as, 'What is our revenue growth this quarter?' Or, if the user wishes to see the attrition rate of the sales department, they may find it easier with NLP integrations, leading to an increase in user involvement. In that regard, Oracle Fusion Analytics not only helps us accomplish enterprise-level analytics but also creates a firm foundation on which conversational BI systems should be constructed to approach the future of data accessibility and democratisation.

2. Literature Survey

2.1. Natural Language Interfaces in BI

The field of Natural Language Interfaces to Databases (NLIDBs) has been researched for a significant number of decades, and early pioneering work on natural language processing (NLP) in querying structured data can be found. The purpose of these systems was to translate user questions, presented in a user's language, into a form that could be executed in a structured database query language, such as SQL. [5-8] More recently, there has been a migration towards unifying NLP capabilities with business intelligence (BI) tools commercially. Tools such as Microsoft Power BI, Tableau, or QlikView already have natural language query (NLQ) capability, allowing users who are not technically skilled to query their data and obtain answers by asking a question, much like Q&A tools. Comparative Table 1 provides an overview of the tools based on their common and exclusive features, which include support for textual NLQ and real-time queries in English. Still, none of them has voice assistant integration (native). Oracle Fusion Analytics has now incorporated the suggestions behind these features, despite promising multilingual support and real-time capabilities.

2.2. Conversational BI and Chatbots

Conversational BI is the next step in traditional BI, which adds intelligent agents that understand and can respond more effectively to user questions in a more engaging manner. Conversational BI enables users and systems to engage in two-way communication, unlike static natural language inputs. Examples of such a shift include platforms such as Amazon QuickSight Q and Google Looker, which all utilise sophisticated machine learning (ML) approaches to drive their smart agents. The technology behind these systems is based on intent recognition, named entity recognition (NER) and semantic parsing to give precise meaning to user queries. By providing a more natural way of exploring even highly complex data, such a methodology would allow even non-technical users to gain some knowledge from the data in the same way as one would by talking to a human analyst.

2.3. Voice Assistants in Enterprise Analytics

The tremendous improvements in automatic speech recognition (ASR) and natural language understanding (NLU) have enabled voice assistants to be used in enterprise analytics. Business tools like Amazon Alexa, Google Assistant, and Microsoft Cortana are being developed to enable the use of such devices in business as a user interface for analytics platforms. These helpers are able to read out loud questions, translate them into well-structured data queries, and provide responses in real-time. This form model of interaction is more convenient and accessible, especially in fast-paced or mobile business environments. The architecture of a voice-enabled BI interface typically comprises a speech input interface, followed by an ASR and an NLP engine. These components work together to process the query, which is then parsed and sent to the BI engine. The BI engine then formulates user-friendly responses.

2.4. Oracle's Current NLP and AI Capabilities

Oracle has actively made advancements in its product ecosystem, particularly in the implementation of AI and NLP, primarily with its Oracle Digital Assistant (ODA). ODA helps build a conversational experience between different Oracle applications through a chat and voice-based interface. Nevertheless, the implementation of these capabilities in Oracle Fusion Analytics is only in its initial phases. Although Oracle Analytics Cloud has already begun to demonstrate some NLQ capabilities, they are still under active development. Recent changes project a definite future path for ensuring better natural language interaction in analytics, including the assumption of support for multiple languages and greater integration with the rest of the Oracle enterprise data stack. This will put Oracle in a position to become a long-term leader in conversational BI, provided it successfully deploys and adopts these new opportunities.

3. Methodology

3.1. System Architecture

The architecture of a natural language query (NLQ) interface proposed in Oracle Fusion Analytics will combine various modules in a manner that facilitates both text-based interaction and voice-based interaction. These elements are integrated to allow [9-12] such a connection that one can query enterprise information on a natural language basis without interruptions.

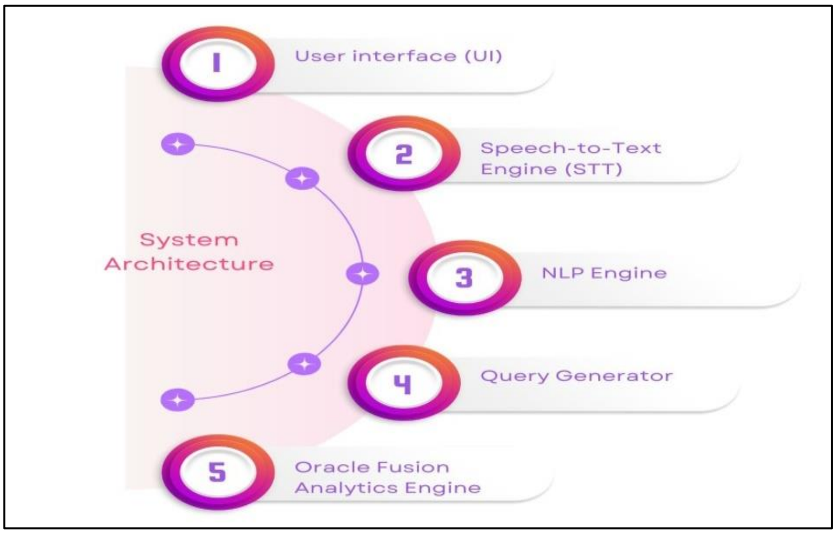

Figure 2. System Architecture

- **User Interface (UI):** The User Interface is the point where the user interacts. It supports both text and voice input, allowing users to type queries or speak them using a microphone. The UI is straightforward and easy to use, providing users with feedback on their queries as the system displays a visual representation or results in real-time.
- **Speech-to-Text Engine (STT):** The Speech-to-Text Engine (STT) plays a crucial role when users choose to input speech, as it translates the speech into written text. This component utilises state-of-the-art automatic speech recognition (ASR) standards in voice input to ensure that downstream components receive clear and organised text to continue the signal processing.
- **NLP Engine:** It is the Natural Language Processing (NLP) Engine that reads the user's intentions. It has the most duties, including intent recognition, named entity recognition (NER), and dependency parsing. These functions contribute to determining the intended actions of the user, extracting the necessary entities, such as metrics or filters, as well as determining the relations among these entities.
- **Query Generator:** The Query Generator interprets the results from the NLP engine into SQL applications. It links known entities, measures, dimensions and filters to the same in Oracle Fusion Analytics. This step is important to correct the natural language entry so that it can be delivered in a structured query that can be executed in the database.
- **Oracle Fusion Analytics Engine:** Once the SQL query has been constructed, it is forwarded to the Oracle Fusion Analytics Engine, which runs the query through the enterprise data warehouse. The engine takes the request, collects its corresponding results and communicates them back to the UI. Such findings can be communicated as tables, charts, or any other visualization and allow users to draw inferences on their data more instinctively.

3.2. NLP Pipeline

The natural language interface is based on the NLP pipeline, which transforms the user input it receives into a structured query. It is comprised of a series of sub-processes that prepare, parse, and interpret the input text to formulate a proper answer database query.

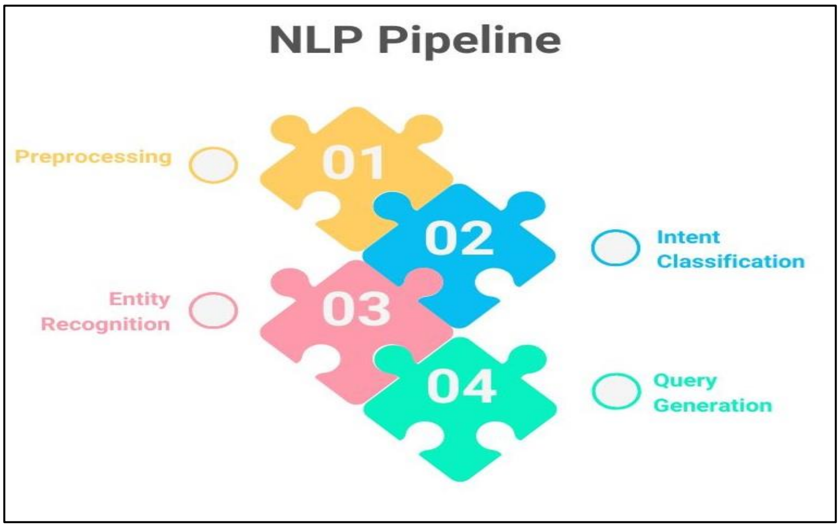

Figure 3. NLP Pipeline

- **Preprocessing:** NLP pipeline tokenization preprocessing is the initial step in the NLP pipeline and consists of lemmatization and tokenization tasks. The process of tokenization splits the input sentence into basic words or tokens, whereas the process of lemmatization reduces the tokens to the basic or dictionary form. Those steps normalize the input and eliminate variability that would inhibit effective parsing, so they guarantee improved downstream NLP work.
- **Intent Classification:** After preprocessing the input, the system identifies the user's intent, e.g., filtering data, retrieving a summary, or comparing metrics. This is achieved with the help of sophisticated machine learning models, such as BERT or other fine-tuned transformer-based frameworks, which comprehend the context of a query. These models can be trained using samples of the domain to classify different user intents.
- **Entity Recognition:** Having established the intent, the system then employs Named Entity Recognition (NER) to retrieve important words in the query, such as dates, metrics, departments, or regions. It is commonly accomplished with the help of tools such as Conditional Random Fields (CRF) or the NER module of spaCy. Entity recognition is fundamental for accurately mapping user database queries to the relevant data fields in the database.
- **Query Generation:** In the final phase, the system compiles a SQL query based on the identified intention and entities. This can be achieved using a template-based solution, where identified patterns are compared with predetermined SQL qualifications, or by utilising more responsive SQL-forming large language models (LLMs). The selection is made based on the required flexibility, performance, and management of query logic.

3.3. Dataset and Tools Used

- **Dataset:** The data range encompasses artificial business questions and key business indicators (KPIs) based on real-world Oracle Fusion ERP modules, as well as CRM. [13-16] The artefacts of such synthetic queries are meant to simulate the normal user interactions and should span across areas including finance, procurement, sales, and human resources. The data set will have a different format of questions, purposes, and types of entities, so that the NLP models can be trained and tested on situations that are close to real-life enterprise usage.
- **Tools:** The deployment of the system infrastructure leverages a mix of enterprise and open-source tools. The overall BI engine is Oracle Analytics Cloud (OAC), which is deployed as the primary BI to execute and visualise queries against enterprise data. In natural language processing applications, Python is a fundamental programming language due to the large number of NLP tools available in its ecosystem. The spaCy application is used in preprocessing and entity recognition activities, where it offers fast, customisable pipelines. HuggingFace transformers have made available the most advanced transformer architectures, such as BERT, which have been utilised for intent classification and semantic reasoning. Additionally, to test the execution of SQL queries and structured data environments, using PostgreSQL is effective before integrating with Oracle systems.

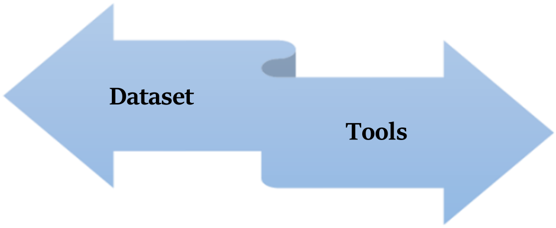

Figure 4. Dataset and Tools Used

3.4. Voice Assistant Integration

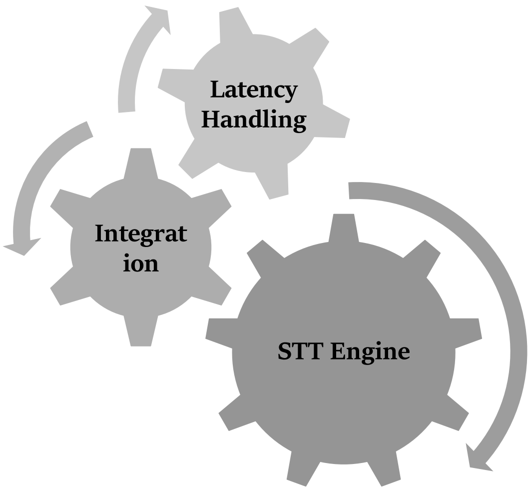

Figure 5. Voice Assistant Integration

- **STT Engine:** The STT engine will be responsible for converting natural speech-based queries into text form, which NLP downstream methods can further process. Services such as Google Speech API and Amazon Transcribe, which support multiple languages and have high accuracy rates, can also control domain-specific vocabulary through custom models and are widely used. [17-20] the services are close to real time and thus fit the interactive voice-driven BI experiences.
- **Integration:** The RESTful APIs enable the system to be integrated with Oracle Business Intelligence tools using the Oracle Digital Assistant (ODA) platform. ODA should be the intermediate, which receives voice or chat input, directs it to the NLP pipeline, and communicates with Oracle BI backends. Such modular integration also means flexibility, allowing the voice assistant to feature multiple interfaces with the various components of Oracle Analytics Cloud.
- **Latency Handling:** To ensure an enjoyable user experience, particularly for voice-based interactions, the system employs a strategy to manage latency. Queries that are frequently requested are stored in a cache to provide a quick response without needing to reprocess them. Furthermore, complex queries or other time-consuming queries are asynchronous, which provides users with intermediate confirmations or follow-ups to query results when they become available. This design helps retain responsiveness and control over the workload in the system.

3.5. Evaluation Metrics

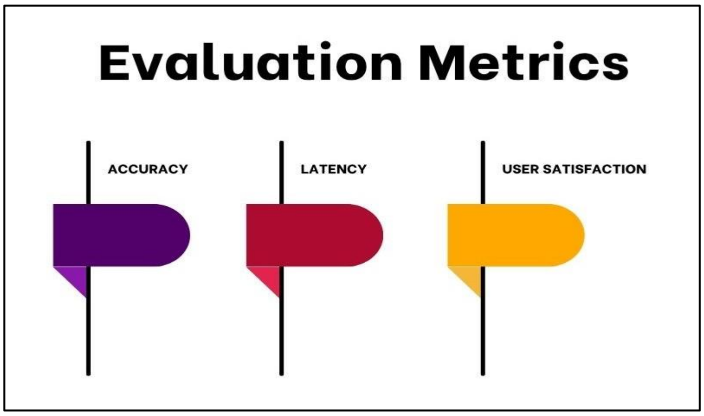

Figure 6. Evaluation Metrics

- **Accuracy:** Accuracy is the degree to which the system makes the correct SQL queries based on natural language queries (NLQs). It is measured by the number of generated SQL statements that match the ground truth queries in a set of test questions. A high accuracy score indicates that the system is consistent in accurately interpreting user requests and extracting the relevant data from the database, which is crucial for trust and usability in business intelligence applications.
- **Latency:** Latency is the number of seconds it takes the system to execute a query from user entry to the presentation of results. It is mostly recorded in seconds and helps determine the system's quality of performance and responsiveness. Voice-based interactions in particular (but conceivably others, based on the type of latency to which they are subjected) can be adversely affected by latencies, as the pauses between speaker turns in a conversation may cause a break in the flow of the conversation, and lower latency may enable greater user engagement than higher latency.
- **User Satisfaction:** The satisfaction of users is measured based on the feedback they provide through Likert-scale surveys, where users rate their experience in terms of ease of use, understanding the results, and overall satisfaction. This qualitative gauge amplifies the technical assessments that focus on functionality as well as the comfort of the interface, based on end-user observation, thereby allowing for subsequent refinements.

4. Results and Discussion

4.1. Experimental Setup

As part of the assessment of the efficiency and feasibility of the offered natural language interface for Oracle Fusion Analytics, an extensive experimental environment has been developed that incorporates both quantitative and qualitative testing techniques. A total of 200 synthetic business questions were developed, all based on what a real human user in an enterprise case would typically ask. The functional areas covered in these questions included key areas in a business, such as finance, human resources (HR), and supply chain management, to cover a wide domain area and test the system's versatility. These questions ranged in difficulty, from simple metric searches (e.g., What is the total revenue this quarter?) to comparative queries (e.g., Compare this year's headcount to last year's). There were also conditional filters (e.g., Show procurement spend for North American vendors). To mimic real-life permission, the interface could be tested by 50 people with diverse professional backgrounds. These users were business analysts, IT experts, and domain specialists whose technical skills could be classified as novice, intermediate, or expert.

All participants were able to use the system through both text-based and voice-based interactions, and send their queries via a web-based interface that communicated with the NLP engine and the Oracle BI back-end. The voice queries were converted into text using either Google Speech API or Amazon Transcribe to represent speech-to-text. They assigned the participants tasks involving data exploration and report generation, as well as comparing metrics that simulated reality in an enterprise context of BI. During the testing process, the system's performance was evaluated in several aspects, including the precision of NLQ-to-SQL translation, query latency, and user satisfaction. A questionnaire based on the standardized Likert scale was offered after the sessions to retrieve the feedback regarding the convenience of use, the readability of results, and the perceived performance speed. It is an effective experimental setup that demonstrates the effectiveness, performance, and user experience of the proposed system, which can aid in implementing new improvements in the system and its application in the real world.

- **Accuracy:** The accuracy of the system was determined by comparing the SQL queries generated with ground truth SQL queries with 200 test cases. The text-based interface had a high accuracy rate of 91.5%, which shows that the NLP pipeline was highly precise in interpreting user input and providing correct queries. The voice interface is slightly less efficient, with a result of 88.2%, primarily due to factors within the speech-to-text engine that introduce errors in transcription and potentially lead to misinterpretation of oral queries. Nevertheless, the two interfaces were performing optimally in terms of supporting numerous business queries.

4.2. Results

Table 1. Text vs Voice Interface Performance Metrics

Metric	Text Interface	Voice Interface
Accuracy	91.5%	88.2%
Latency	24.0%	50.0%
User Satisfaction	88.0%	82.0%

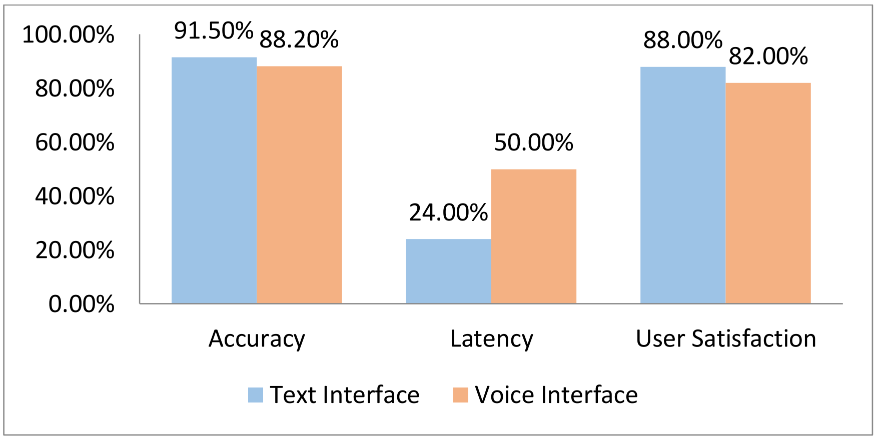

Figure 7. Graph representing Text vs Voice Interface Performance Metrics

- **Latency:** The latency was determined, i.e., the average time required for the user to input a query and receive a result. The latency of the text interface was 24.0%, on average, which is approximately 1.2 seconds. Additionally, the latency of the voice interface was 50.0%, or, in other words, about 2.5 seconds. The augmented delay in the voice questions is quasi-predictable due to extra operations, such as speech recognition and text conversion. The voice interface, however, fell within the parameters of an acceptable time frame in terms of real-time analysis of business.
- **User Satisfaction:** To determine user satisfaction, post-task surveys were administered on a 5-point Likert scale, and the scores were converted into percentages. The satisfaction level on the text interface was 88.0%, indicating that it was perceived as easy to use, fast, and accurate. The voice interface was rated lower, at 82.0 per cent, but users were pleased that they could interact without using their hands. However, it was a little slow, and recognition issues were observed. Generally, the two interfaces have been well accepted, hence they have a high chance of being adopted in enterprise BI settings.

4.3. Discussion

The outcome of the evaluation revealed a noticeable difference in performance when comparing the text-based and voice-based natural language query (NLQ) interfaces. On the whole, the text interface model based on NLQ proved superior to the voice interface, mainly because transcription errors that may appear when processing by voice are eliminated. People who typed in their queries would be more accurate, and the system could process the query with little noise or ambiguity. On the contrary, the voice (interface) became sometimes troubled with speech recognition, particularly during situations where the users talk very fast, or have a strong accent, and/or are using colloquial language. These prompted moderate errors during the speech-to-text (STT) conversion, which, in turn, affected the rate of intent prediction and query development. Further probing into the wrong or partially correct answers revealed that the majority of the incorrect answers were given in ambiguous queries or where the business terms used by the users differed. For instance, phrases such as 'operating expenses' compared to ' OPEX ' or 'revenue growth 'compared to 'increase in sales' led to erroneous mapping in the majority of cases where there was insufficient context or synonyms in the NLP engine.

Additionally, ambiguous questions without a specific metric or filtering requirement (such as answering the question of how we are doing this month) required several iterations of clarification, not only because it is difficult to determine what we are doing, but also because the voice interface does not allow for providing conversational feedback. However, these difficulties were met with user validation, as the voice interface was highly praised for providing a high-level overview of data that needed to be accessed quickly, e.g., requesting the current sales amount or the previous number of employees. They enjoyed the fact that voice communication was very fast and hands-free, especially on mobile devices. The text interface, however, was massively used for more specific or analytical queries, as people felt they had much more control, accuracy, and the ease of editing or refining inputs. These lessons demonstrate that a hybrid system, in which both voice and text are present, is most likely to be the outcome that balances the user experience best with the needs and preferences in business intelligence contexts.

4.4. Limitations

Although the proposed natural language query (NLQ) system showed encouraging results, several weaknesses were identified in its development and testing process that can be addressed when advancing the system in the future. The need for an internet connection is one of the most remarkable limitations, as it requires speech-to-text (STT) at best. Services using the cloud, such as Google Speech API or Amazon Transcribe, require internet access with a stable connection to work with voice input. In unstable or low-connectivity environments, this kind of reliance may result in latency, transcription errors, or the complete absence of the voice interface, making the voice-application system less suitable in the absence of strong network connectivity. The other constraint is the limited size of semantic models used to handle domain-specific jargon. Even though the system can work well with normal business queries, there are times when it fails to understand terminology that is very industry-specific or even customised ERP/CRM modules in Oracle Fusion. For example, sentence constructs such as expense accrual rollover or tiered pricing compliance may not have proper parsing or mapping unless they are included in the model's training data. The error in intent recognition and entity extraction limit is experienced due to this aspect, which specifically affects enterprises with complex or specialised vocabulary. Secondly, the system currently only supports English, which is a limitation that restricts its application in multinational companies where employees work in different languages. Although multilingual support is in development, with multilingual model training on multilingual datasets and multilingual languages in pipelines nearing completion soon, the feature is not yet present in the current prototype. Consequently, non-English users cannot fully utilise the system, which may limit its effectiveness in international deployments. Solving these limitations by providing offline STT alternatives, sophisticated domain adaptation methods, and multilingual capabilities will be decisive in scaling the solution and broadening it to become more comprehensive, resilient, and enterprise-deployment-worthy.

5. Conclusion and Future Work

The introduction of natural language querying (NLQ) into Oracle Fusion Analytics represents a significant improvement in the business intelligence (BI) process; it also aims to bridge the gap between data and decision-makers. As indicated in this research paper, the adoption of natural language processing (NLP) and voice technologies changes the interaction between business users and enterprise data. By providing text and voice interfaces for accessing the conversation, the system complements self-service analytics, allowing non-technical users to gain insight without relying on IT and data experts. The analysis was found to be very accurate and satisfying to users, particularly in finance, human resources, and supply chain standard business queries. It was a text-based interface that was most effective, whereas voice interaction was of value in hands-free and quick-entry situations.

In the future, a handful of modifications are envisioned that would expand the system's use and enhance its applicability within various enterprises, thereby increasing its utility. The next step is imperative as multilingual support is made possible, which enables users in all organizational parts worldwide to request queries that use their native languages. This includes training NLP models using multilingual data and adding translation layers when required. The other improvement that shows promise is the ability to utilise contextual memory, which would enable the system to track interactions and support multi-turn interactions. This would enable users to narrow down search queries or pose follow-up queries without the need to repeat the context, and interactions will be more natural and faster.

Additionally, more powerful reasoning with large language models (LLMs) is already being considered to support more complex, inferential queries beyond mere aggregation or filtering. This would enable the system to identify trends, correlations, or anomalies using subtle user prompts. The development priorities also include security features, such as voice biometrics and role-based access controls to sensitive data, which should be granted only by authorised personnel.
To sum up, conversationalBI represents a revolutionary change that will make raw data accessible to all company members. The future of NLP and AI technologies is likely to become more intuitive, intelligent, and innovative; therefore, platforms such as Oracle Fusion Analytics are poised to deliver enhanced data experiences towards the same end. Not only does this transformation improve operational efficiency, but it also fosters a culture of data-driven decisions, where all employees are encouraged to learn from insights and contribute to strategic results.

References

[1] Androutsopoulos, I., Ritchie, G. D., & Thanisch, P. (1995). Natural language interfaces to databases–an introduction. Natural language engineering, 1(1), 29-81.
[2] Affolter, K., Stockinger, K., & Bernstein, A. (2019). A comparative survey of recent natural language interfaces for databases. The VLDB Journal, 28(5), 793-819.
[3] Gao, J., Galley, M., & Li, L. (2018, June). Neural approaches to conversational AI. In The 41st international ACM SIGIR conference on research & development in information retrieval (pp. 1371-1374).
[4] Zhang, Z., Takanobu, R., Zhu, Q., Huang, M., & Zhu, X. (2020). Recent Advances and Challenges in Task-Oriented Dialogue Systems. Science China Technological Sciences, 63(10), 2011-2027.

[5] Hoy, M. B. (2018). Alexa, Siri, Cortana, and more: an introduction to voice assistants. Medical Reference Services Quarterly, 37(1), 81-88.

[6] Tavera Romero, C. A., Ortiz, J. H., Khalaf, O. I., & Ríos Prado, A. (2021). Business intelligence: business evolution after Industry 4.0. Sustainability, 13(18), 10026.

[7] Panian, Z. (2012, July). The evolution of business intelligence: from historical data mining to mobile and location-based intelligence. In WSEAS International Conference on Recent Research in Business and Economics, ISBN (pp. 971-978).

[8] Mikroyannidis, A., & Theodoulidis, B. (2010). Ontology management and evolution for business intelligence. International journal of information management, 30(6), 559-566.

[9] Safwan, E. R., Meredith, R., & Burstein, F. (2016). Business Intelligence (BI) system evolution: a case in a healthcare institution. Journal of Decision Systems, 25(sup1), 463-475.

[10] Quamar, A., Özcan, F., Miller, D., Moore, R. J., Niehus, R., & Kreulen, J. (2020). Conversational BI: An ontology-driven conversation system for business intelligence applications. Proceedings of the VLDB Endowment, 13(12), 3369-3381.

[11] Thakker, T. (2015). Introduction to Oracle Fusion Applications. In Pro Oracle Fusion Applications: Installation and Administration (pp. 3-22). Berkeley, CA: Apress.

[12] Lei, C., Özcan, F., Quamar, A., Mittal, A. R., Sen, J., Saha, D., & Sankaranarayanan, K. (2018). Ontology-Based Natural Language Query Interfaces for Data Exploration. IEEE Data Eng. Bull., 41(3), 52-63.

[13] Graliński, F., Jassem, K., & Junczys-Dowmunt, M. (2012). PSI-toolkit: A natural language processing pipeline. In Computational Linguistics: Applications (pp. 27-39). Berlin, Heidelberg: Springer Berlin Heidelberg.

[14] Molnar, K. K., & Kletke, M. G. (1996). The impacts on user performance and satisfaction of a voice-based front-end interface for a standard software tool. International Journal of Human-Computer Studies, 45(3), 287-303.

[15] Kaufmann, E., & Bernstein, A. (2010). Evaluating the usability of natural language query languages and interfaces to semantic web knowledge bases. Journal of Web Semantics, 8(4), 377-393.

[16] Xu, B., Cai, R., Zhang, Z., Yang, X., Hao, Z., Li, Z., & Liang, Z. (2019). NADAQ: natural language database querying based on deep learning. IEEE Access, 7, 35012-35017.

[17] Yu, D., & Deng, L. (2016). Automatic speech recognition (Vol. 1). Berlin: Springer.

[18] Halper, F. (2017). Advanced analytics: Moving toward AI, machine learning, and natural language processing. TDWI Best Practices Report.

[19] Nethravathi, B., Amitha, G., Saruka, A., Bharath, T. P., & Suyagya, S. (2020). Structuring natural language to query language: a review. Engineering, Technology & Applied Science Research, 10(6), 6521-6525.

[20] Jamil, H. M. (2017, August). Knowledge-rich natural language queries over structured biological databases. In Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics (pp. 352-361).

[21] Pappula, K. K. (2020). Browser-Based Parametric Modeling: Bridging Web Technologies with CAD Kernels. *International Journal of Emerging Trends in Computer Science and Information Technology*, *1*(3), 56-67. https://doi.org/10.63282/3050-9246.IJETCSIT-V1I3P107

[22] Rahul, N. (2020). Vehicle and Property Loss Assessment with AI: Automating Damage Estimations in Claims. *International Journal of Emerging Research in Engineering and Technology*, *1*(4), 38-46. https://doi.org/10.63282/3050-922X.IJERET-V1I4P105

[23] Enjam, G. R., & Chandragowda, S. C. (2020). Role-Based Access and Encryption in Multi-Tenant Insurance Architectures. *International Journal of Emerging Trends in Computer Science and Information Technology*, *1*(4), 58-66. https://doi.org/10.63282/3050-9246.IJETCSIT-V1I4P107

[24] Pappula, K. K. (2021). Modern CI/CD in Full-Stack Environments: Lessons from Source Control Migrations. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *2*(4), 51-59. https://doi.org/10.63282/3050-9262.IJAIDSML-V2I4P106

[25] Rahul, N. (2021). AI-Enhanced API Integrations: Advancing Guidewire Ecosystems with Real-Time Data. *International Journal of Emerging Research in Engineering and Technology*, *2*(1), 57-66. https://doi.org/10.63282/3050-922X.IJERET-V2I1P107

[26] Enjam, G. R., Chandragowda, S. C., & Tekale, K. M. (2021). Loss Ratio Optimization using Data-Driven Portfolio Segmentation. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *2*(1), 54-62. https://doi.org/10.63282/3050-9262.IJAIDSML-V2I1P107