

Complex Business Intelligence Queries in Natural Language

Diptiben Ghelani¹

¹Department of Computer engineering, Gujrat Technological University

September 23, 2022

Complex Business Intelligence Queries in Natural Language

Diptiben Ghelani

Department of Computer engineering, Gujrat Technological University, Ahmedabad, Gujrat, India

ABSTRACT: Natural Language Interface to Database (NLIDB) eliminates the need for an end user to use complex query languages like SQL by translating the input natural language state-ments to SQL automatically. Although NLIDB systems have seen rapid growth of interest recently, the current state-of-the-art systems can at best handle point queries to retrieve certain column values satisfying some filters, or aggrega-tion queries involving basic SQL aggregation functions. In this demo, we showcase our NLIDB system with extended capabilities for business applications that require complex nested SQL queries without prior training or feedback from human in-the-loop. In particular, our system uses novel al-gorithms that combine linguistic analysis with deep domain reasoning for solving core challenges in handling nested queries. To demonstrate the capabilities, we propose a new benchmark dataset containing realistic business intelligence queries, conforming to an ontology derived from FIBO and FRO financial ontologies. In this demo, we will showcase a wide range of complex business intelligence queries against our benchmark dataset, with increasing level of complexity. The users will be able to examine the SQL queries generated, and also will be provided with an English description of the interpretation.

INTRODUCTION

With the omnipresence of mobile devices coupled with recent advances in automatic speech recognition capabilities, there has been a growing demand for Natural Language Interfaces to Databases (NLIDB). The reason behind the rapid increase of popularity of NLIDB systems is due to the fact that they do not require the users to learn a complex query language, such as SQL, or understand the exact schema of the data, or how it is stored, making it very easy to explore complex data sets, beyond simple keyword search queries. Natural language interfaces to databases have become an increasingly important research field during the last decade [6, 7, 9, 10]. Most prominent among state-of-the-art includes our previous work ATHENA [6, 10] and its predecessor NALIR [7], both of which essentially are rule-based systems. While NALIR supports few specific types of basic nested queries using natural language parsing and human in-the-loop feedback, it could not handle complex nested queries, and the need of an expert feedback became a bottleneck. On the other hand, ATHENA improves over NALIR by avoiding user dependence with ontologies capturing domain seman-tics, but it falls short of handling nested queries as well. A different category of works [11, 13] employs machine learning approaches to tackle the problem. However, the main drawback with this approach is the limited amount of domain-specific data to train with, making systems like Seq2SQL [13] to handle only single table select and project

queries. More recently, DBPal [12] claims to handle join queries and a selective set of nested queries where the corresponding query phrasings have to be available as part of the training data.

In both segments of research, nested query handling remains an open challenge, because none of the systems are equipped with the deep domain reasoning capability, which is essential for interpreting BI queries and producing correct nested queries. In particular, the challenges in handling nested queries are as follows.

Detecting a NL query which requires nesting. Whether a natural language (NL) query requires nesting depends on multiple convoluted factors, including operations involved, operands in use, domain semantics, and SQL query semantics. Forming the correct subqueries. To form the right subqueries (inner and outer), we need to accurately identify which tokens from the NL query are to be considered for each subquery, which is often non-trivial in natural language phrasing. Forming the correct join condition between sub-queries. Finding the correct join conditions between sub-queries depends not only on SQL query semantics, but often requires deep domain reasoning. In this demonstration, we will showcase our NLIDB system, which is an extension of our earlier work on ATHENA. The system is capable of handling a wide range of BI queries including complex join and nested queries. To address the above challenges, our system uses novel algorithms that combine linguistic analysis with deep domain reasoning without prior training or feedback from human in-the-loop. In our previous work [5], we demonstrated how to use ATHENA in querying financial knowledge bases curated from unstructured content. In this demonstration, we enrich the Finance domain ontology beyond [5], with transaction data, resulting in a new benchmark FIBEN with more concepts and relations. This new benchmark emulates data marts, and hence allows more complex BI queries. In particular, we integrate two widely known financial domains (SEC [3] and TPoX [8]) to create FIBEN, and curate meaningful data conforming to it. We demonstrate the capabilities of our system with real life BI queries and complex analytic queries on FIBEN to derive crucial business insights from underlying data, with increasing level of complexity. The users will be able to examine the SQL queries generated, and also will be provided with an English description of the interpretation¹.

2 A FINANCE BENCHMARK DATASET

FIBEN is a complex business intelligence benchmark dataset that emulates a real-world data mart. To create FIBEN, we combined data from two different financial sub-domains. SEC dataset [3] is a publicly available dataset extracted from public Securities and Exchange Commission (SEC) filings submitted as XBRL documents. The filing data is curated using an enrichment flow [5], and provides information

about public companies, their officers, and financial metrics reported over a period of time. This dataset is a source of archived financial data reported to the SEC over the last 15 years, and lends itself to provide rich business insights regarding the financial health and performance of public companies in a variety of different industry sectors.

TPoX dataset [8] is a transaction processing bench-mark, which describes financial transactions over holdings and securities provided by public companies. We used the TPoX data generator to generate dynamic data in terms of financial transactions. Each financial transaction is linked to a customer account. The customer account describes the customer’s portfolio in terms of the securities held, and is associated with the buying or selling of securities such as stocks, bonds and mutual funds of publicly traded companies in different financial markets over a period of time. We mapped the two datasets through an extensive set of data transformations to the relevant subsets of two standard finance ontologies: Finance Industry Business Ontology (FIBO) [1] and Finance Report Ontology (FRO) [2]. FIBO is the de-facto industry standard defined by the enterprise data management (EDM) council to represent business concepts and information in the finance domain. FRO is a formal report ontology of an XBRL based financial report which captures the financial metric data reported by public companies to SEC. Figure 1 shows the combined FIBO-FRO ontology describing the schema of our FIBEN dataset.

To provide a standard benchmark for testing BI applications, we integrate the two datasets through public companies. Securities held and traded by customers in the TPoX dataset are provided by public companies that are available in the SEC dataset.

SYSTEM OVERVIEW

The system architecture is depicted in Figure 3 and is extended from the architecture presented in ATHENA [10]. The components specific to nested query handling include Nested Query Detector, Subquery Formation, and Subquery Join Condition. While the complete details of each of the components is beyond the scope of a demonstration paper, we will use an example query flow to explain these extended components. Consider the query “who bought more IBM stocks than they sold”.

Evidence Annotator. This follows the same process as detailed in ATHENA [10], where the Evidence Annotator scans through all the words (a.k.a. tokens) and identifies “stock” to be a mention of the ontology concept “ListedSecurity”, “IBM” as a value in the database column corresponding to “ListedSecurity.name”. Similarly, “bought” and “sold” are values of “Transaction.type”, and “who” is an indication to include “Persons” in the select clause of the target SQL query. The tokens that are mapped to some ontology elements are called entities.

of “more than” operation must be expanded into a subquery in order to be compared using a numeric operator like “more than”.

Subquery Formation. Once a query is detected to be a nested query, subquery formation module builds two sub-queries for the inner and outer queries, by identifying the right set of tokens associated with each part. Note that the subqueries do not need to have a disjoint set of tokens and often the subqueries share

tokens without any hard boundary on their positions in the NL query. The right set of tokens for each subquery is found by using a set of rules applicable to the query, where each rule considers the annotator outputs and domain elements. In the example query, the right set of tokens for the outer query are {Who, Bought, IBM, Stock}, and for inner query are { Who, Sold, IBM, Stock}, where “Who”, “IBM”, “Stock” tokens are shared between both sub-queries. The set of rules applied to determine the shared tokens for this specific question are as follows.

Rule 1. Any argument(“Stock”) of a numeric comparison (“more than”) will be shared across subqueries.

Rule 2. Any instance value (“IBM”) from the database having a functional relation with an argument (“Stock”) of a comparison argument will be shared.

Rule 3. An entity in the outer Query (“who”) that is co-referenced in the inner query (“they”) will be shared with the inner query as well.

Subquery Join Condition. Once the logical subqueries are formed, the join condition needs to be identified between the subqueries to produce the complete query. The join condition depends on the linguistic analysis as well as the domain reasoning. In the example question, linguistic analysis maps “more than” to “>” operator, and identifies that it compares with subquery results. Domain analysis finds the argument “stock” is not non-numeric, and hence it cannot be the operand of the comparison. It is then left to domain reasoning to identify that every stock has an associated “count” for each transaction and “count” is a numeric entity. Join condition is thus derived as “>” applied over “SUM (Transaction.hasCount)”.

4 DEMONSTRATIONS

As mentioned earlier, our work focus on handling a wide range of complex nested queries without the need for human-

Nested Query Detector. Multiple aspects of a natural language query along with the domain semantics can indicate if it requires nesting. The different aspects include linguistic analysis of the NL query and semantic annotators which maps specific tokens from NL query to domain terms. There is a reasoning submodule that works on the outputs of linguistic analyzer and semantic annotators to detect a possible nested query. In this example, linguistic analysis identifies a comparison of type “more than” with argument “stocks” and Semantic Annotators maps “stock” to a concept. The reasoning submodule identifies that the operand “stocks”

in-the-loop feedback or prior training. We will demonstrate this through three different tasks.

Familiarization with FIBEN Ontology. The complete FIBEN ontology consists of 152 concepts, 664 data properties and 159 object properties (i.e., relations). Such large-scale on-tology emphasizes why an end user needs to get familiarized with the domain semantics in order to appreciate the insight-ful analytic queries in this domain. We plan to showcase the complete ontology and the underlying database schema to the audiences so that they understand the various actors in the finance domain and the different roles that those ac-tors are playing. They can use tools like WebVOWL [4] to navigate the ontology themselves. We will focus on the wide range of complex nested queries that our system can handle. Instead of a story script, we would show a sequence of queries with increasing order of their complexity, and pre-populate the UI to have such sample complex business intelligence queries. An example snapshot of such a UI screen is shown in Figure 4. As can be seen, the queries cover a wide range of complexity, ranging from queries that do not need nesting and can be expressed using having clauses, to queries with different join conditions combining inner and outer query blocks in the nested query. For example, the query “Show me everyone who bought and sold the same stock” needs equality check (for “same”) among non-numeric fields. Whereas a query like “Who has bought more IBM stocks than they sold” is even more com-plex because it involves numeric comparison between two aggregation (i.e., sum of stock counts) results. **Interactive UI.** The same UI with pre-populated interest-ing business intelligence queries (Figure 4) will be used to let the audiences interact with the system. The audiences can choose to run such pre-populated queries, or frame their own queries over the FIBEN ontology for an unscripted demon-stration. Figure 5 shows the visualization with an example question and its answer. For every query, we will also show the English description of how the system interpreted the user statement, and the generated SQL query which is ex-ecuted against the data store to produce the query result.

Our demo will run on DB2®². Figure 6 shows the part of the ontology graph which was relevant for answering the query “Show me each transaction for IBM whose price is less than the average selling price”. It also shows the specific subgraphs which were created for both outer query and in-ner query. We will include these details for the audiences to easily correlate the query interpretation with the underlying domain semantics.

REFERENCES

- [1] 2018. FIBO. <https://spec.edmcouncil.org/fibo/>. (Oct. 2018).
- [2] 2018. FRO. <http://xbrl.squarespace.com/financial-report-ontology/>. (Oct. 2018).
- [3] 2019. SEC Financial Statement Data Sets. <https://www.sec.gov/dera/data/financial-statement-data-sets.html>. (Jan. 2019).
- [4] 2019. WebVOWL: Web-based Visualization of Ontologies. <http://vowl.visualdataweb.org/webvowl.html>. (Jan. 2019).
- [5] Shreyas Bharadwaj, Laura Chiticariu, Marina Danilevsky, et al. 2017. Creation and Interaction with Large-scale Domain-Specific Knowledge Bases. PVLDB 10 (2017), 1965–1968.
- [6] Hepp, M., et al., *Ontology management: semantic web, semantic web services, and business applications*. 2007: Springer Science & Business Media.
- [7] Lal, M., *Web 3.0 in Education & Research*. BVICAM's International Journal of Information Technology, 2011. 3 (2).
- [8] Miranda, P., P. Isaias, and C. J. Costa, *E-Learning and web generations: Towards Web 3.0 and E-Learning 3.0*. International Proceedings of Economics Development and Research, 2014. 81: p. 92.
- [9] Garrigos-Simon, F. J., R. L. Alcamí, and T. B. Ribera, *Social networks and Web 3.0: their impact on the management and marketing of organizations*. Management Decision, 2012.
- [10] Tasner, M., *Marketing in the moment: the practical guide to using Web 3.0 marketing to reach your customers first*. 2010: Ft Press.
- [11] Erragcha, N. and R. Romdhane, *New faces of marketing in the era of the web: from marketing 1.0 to marketing 3.0*. Journal of research in marketing, 2014. 2 (2): p. 137-142.
- [12] Ferrari, S., *Marketing strategies in the age of web 3.0*, in *Mobile Computing and Wireless Networks: Concepts, Methodologies, Tools, and Applications*. 2016, IGI Global. p. 2132-2149.
- [13] Polat, V. and A. E. Akgün, *A conceptual framework for marketing strategies in web 3.0 age: adaptive marketing capabilities*. Journal of Business Studies Quarterly, 2015. 7 (1): p. 1.
- [14] Poore, M., *The Next G Web. Discernment, meaning-making, and the implications of Web 3.0 for education*. Technology, pedagogy and education, 2014. 23 (2): p. 167-180.
- [15] Horrocks, I., et al. *Semantic web architecture: Stack or two towers?* in *International Workshop on*

Principles and Practice of Semantic Web Reasoning. 2005. Springer.

- [16]Rudman, R., *Web 3.0: governance, risks and safeguards*. Journal of Applied Business Research (JABR), 2015. 31 (3): p. 1037-1056.
- [17]Tiago, M. T. P. M. B. and J. M. C. Veríssimo, *Digital marketing and social media: Why bother?* Business horizons, 2014. 57 (6): p. 703-708.
- [18]Glazier, A., *Searchial Marketing:: How Social Media Drives Search Optimization in Web 3.0*. 2011: Author House.
- [19]Hatzivasilis, G., et al. *Secure Semantic Interoperability for IoT Applications with Linked Data*. in *2019 IEEE Global Communications Conference (GLOBECOM)*. 2019. IEEE.
- [20]Ciccarese, P., et al. *An open annotation ontology for science on web 3.0*. in *Journal of biomedical semantics*. 2011. BioMed Central.
- [21]Bergmann, H., et al., *Semantic interoperability to enable smart, grid-interactive efficient buildings*. 2020, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States).
- [22] Patela, R. and N. Patela, *The Third Generation of Internet: The Semantic Web as a Component of Web 3.0*.
- [23]Parida, V., M. Westerberg, and J. Frishammar, *Inbound open innovation activities in high-tech SMEs: the impact on innovation performance*. Journal of small business management, 2012. 50 (2): p. 283-309.
- [24]Lindberg, K., *Economic impacts*. The encyclopedia of ecotourism, 2001. 5 (23): p. 363-377.
- [25]Piazolo, M. and N. A. Zanca, *Medical tourism: A case study for the USA and India, Germany and Hungary*. Acta Polytechnica Hungarica, 2011. 8 (1): p. 137-160.
- [26]Kherfi, M. L., D. Ziou, and A. Bernardi, *Image retrieval from the world wide web: Issues, techniques, and systems*. ACM Computing Surveys (Csur), 2004. 36 (1): p. 35-67.
- [27]Cambria, E. and A. Hussain, *Sentic computing*. marketing, 2012. 59 (2): p. 557-577.
- [28]Knautz, K., et al., *Finding emotional-laden resources on the World Wide Web*. Information, 2011. 2 (1): p. 217-246.
- [29]Allemang, D. and J. Hendler, *Semantic web for the working ontologist: effective modeling in RDFS and OWL*. 2011: Elsevier.
- [30]Kiryakov, A., et al., *Semantic annotation, indexing, and retrieval*. Journal of Web Semantics, 2004. 2 (1): p. 49-79.

- [31] Russo, A. and F. Perrini, *Investigating stakeholder theory and social capital: CSR in large firms and SMEs*. Journal of Business ethics, 2010. 91 (2): p. 207-221.
- [32] Ghelani, D., & Hua, T. K. (2022). Conceptual Framework of Web 3.0 and Impact on Marketing, Artificial Intelligence, and Blockchain. *International Journal of Information and Communication Sciences*, 7(1), 10.
- [33] Ghelani, D., & Hua, T. K. *A Perspective Review on Online Food Shop Management System and Impacts on Business*.
- [34] Ghelani, D., Hua, T. K., & Koduru, S. K. R. (2022). A Model-Driven Approach for Online Banking Application Using AngularJS Framework. *American Journal of Information Science and Technology*, 6(3), 52-63.
- [35] Dr. John Ughulu. The role of Artificial intelligence (AI) in Starting, automating and scaling businesses for Entrepreneurs.. *ScienceOpen Preprints*. DOI: 10.14293/S2199-1006.1.SOR-PP5ZKWJ.v1
- [36] Ughulu, J. *Entrepreneurship as a Major Driver of Wealth Creation*.
- [37] Oak, R., Du, M., Yan, D., Takawale, H., & Amit, I. (2019, November). Malware detection on highly imbalanced data through sequence modeling. In *Proceedings of the 12th ACM Workshop on artificial intelligence and security* (pp. 37-48).
- [38] Hua, T. K., & Biruk, V. (2021). *Cybersecurity as a Fishing Game: Developing Cybersecurity in the Form of Fishing Game and What Top Management Should Understand*. Partridge Publishing Singapore.