

## AI-Augmented LangChain: Facilitating Natural Language SQL Queries for Non-Technical Users

Arpan Shaileshbhai Korat

School of Engineering and Applied Sciences, State University of New York - Buffalo, NY, USA

### ABSTRACT

This study explores the potential of LangChain, a framework for constructing applications with advanced language models, to translate natural language queries into executable SQL code. Study propose an innovative LangChain-based architecture that receives a natural language query, analyzes it with a language model, and generates the corresponding SQL statement for database querying. This approach aims to empower non-technical users, facilitate inter-team collaboration, and enable data-informed decision-making. However, challenges persist, including managing complex queries and grasping domain-specific terminology. This research investigates the methodology and system design of our proposed natural language interface for databases, leveraging LangChain and extensive language models. Study also explore the possibilities and potential applications of this system, as well as future research avenues for enhancing its functionalities and addressing current constraints. By integrating advanced natural language processing with database technologies, this research aims to enable inclusive and powerful data querying experiences.

### \*Corresponding author

Arpan Shaileshbhai Korat, School of Engineering and Applied Sciences, State University of New York - Buffalo, NY, USA.

**Received:** June 17, 2024; **Accepted:** June 20, 2024, **Published:** June 28, 2024

**Keywords:** LangChain, Natural Language Processing (NLP), SQL Translation, Natural Language Interface, Database Querying, Data-Informed Decision-Making

### Introduction

Discover a world where interacting with databases is as simple as having a chat. Enter the domain of Natural Language to SQL (NL2SQL), where LangChain's state-of-the-art technology bridges the divide between user-friendly queries and intricate SQL commands. This detailed manual will guide us through an immersive experience to unleash the full potential of data, ensuring accessibility for all, regardless of technical know-how. From healthcare to finance, the possibilities of NL2SQL are endless, enabling businesses to harness the power of their data through intuitive natural language interactions. In this article, delve into a step-by-step process of creating an intelligent NL2SQL model using LangChain, covering everything from building a foundational model to integrating advanced features such as few-shot learning, dynamic table selection, prompt customization, and conversational memory. By the end of this in-depth exploration, all of us have the expertise and resources to develop an NL2SQL system that not only converts natural language queries into SQL commands but also retains context, personalizes responses, and delivers a seamless user experience, democratizing data access through- out your organization.

As embark on this exciting journey, develop a deep understanding of the following:

- **Building a Solid NL2SQL Model:** Establishing the foundation for converting natural language queries into SQL commands.

- **Utilizing Few-Shot Learning:** Improving model performance by strategically incorporating examples.
- **Implementing Dynamic Few-Shot Example Selection:** Adapting example selections to match the query context and ensure relevance.
- **Enabling Dynamic Relevant Table Selection:** Empowering the model to automatically identify the appropriate tables to query based on the natural language input.
- **Customizing Prompts and Responses:** Fine-tuning the model's interactions to provide clear, concise, and relevant responses.
- **Integrating Conversational Memory:** Equipping the model with the ability to handle follow-up questions by retaining the conversation's context.

Throughout each stage, this research delves into the underlying concepts, providing a comprehensive guide to the implementation process and demonstrating the tangible out-comes. By the conclusion, the necessary tools and knowledge to fully leverage the power of NL2SQL for databases will be provided. The goal is to make querying databases as intuitive as conversing in natural language, thereby unlocking the true potential of data insights. This study aims to transform the approach to database management, emphasizing the seamless integration of natural language processing techniques to enhance data accessibility and utilization. The outcomes of this research will not only facilitate more efficient data querying but also contribute to the broader field of database management and natural language processing.

### Literature Survey

The challenge of translating natural language queries into SQL has

been an active area of research for several decades, with various approaches proposed and explored.

### **Rule-Based Approaches in the Traditional Sense**

Early research on translating natural language to SQL focused on the use of rule-based systems. These systems utilized carefully crafted rules, grammars, and lexicons to parse natural language queries and convert them into corresponding SQL components. However, these systems heavily relied on manual effort to define the rules and encode domain-specific knowledge.

One of the pioneering systems in this field was PRECISE (Programming Rachelle in a Controlled English), which was developed in the late 1970s. PRECISE used a controlled subset of English and a set of syntax rules to translate queries into an intermediate logical representation, which was then converted into SQL. Although PRECISE demonstrated the potential of rule-based approaches for natural language to SQL translation, it was limited to a specific domain and query complexity.

Another notable system from this era was NALIR (Natural Language Interface for Relational databases), developed in the 1980s. NALIR employed a sophisticated grammar and lexicon to parse natural language queries and generate an intermediate semantic representation. This representation was then mapped to SQL using a set of translation rules. While NALIR achieved reasonable accuracy for various query types, it was hindered by the need for manual effort in defining and maintaining the rules and domain knowledge.

Although these rule-based systems showed promise in constrained domains, they faced significant challenges when it came to scaling and adapting to new domains or query complexities. The labor-intensive nature of defining rules and encoding domain knowledge made these systems difficult to maintain, which limited their widespread adoption.

### **Statistical and Machine Learning Approaches**

Advancements in machine learning techniques led researchers to explore the use of statistical models and neural networks to directly analyze the connections between natural language and SQL from data.

An early adopter of this approach was NL2SQL, which employed a sequence-to-sequence model with an encoder-decoder architecture [1]. The encoder processed the natural language query, while the decoder generated the corresponding SQL statement. This system demonstrated the ability to handle more complex queries compared to rule-based systems and showed superior generalization capabilities. However, it required a significant parallel dataset of natural language queries and their corresponding SQL statements for training, which could be costly and time-consuming to obtain.

Building on this work, SQLNet introduced a more sophisticated neural network architecture that combined a sequence-to-set model with a sketch-based method [2]. This system initially predicted a high-level sketch of the SQL query, which was then refined and completed with specific components. SQLNet achieved excellent performance on various benchmark datasets but still relied on extensive parallel datasets for training.

While these machine learning-based approaches showed promising results and improved generalization capabilities, they faced challenges in handling complex queries and domain-specific

language. Additionally, the need for large parallel datasets limited their applicability to domains or query types where such data was not readily available or expensive to acquire.

### **Semantic Parsing and Intermediate Representations**

Another area of research focused on the conversion of natural language queries into intermediate logical representations, which could then be transformed into SQL. This particular approach aimed to separate the understanding of natural language and the generation of SQL, potentially allowing for better adaptability and generalization to new query types or domains.

The NALIR system, mentioned previously, utilized a semantic parsing technique to translate natural language queries into an intermediate representation known as Query Representation Language (QRL) [2]. This representation captured the semantic structure of the query, which was then mapped to SQL using a set of predefined rules.

In a similar vein, Zhang et al. proposed an approach that employed a semantic parser to convert natural language queries into an Abstract Meaning Representation (AMR) [3]. This intermediate representation also captured the semantic structure of the query, which was then converted to SQL using a combination of rules and domain-specific knowledge.

These approaches showcased the potential of utilizing intermediate representations to enhance generalization and adaptability. However, they still relied on rule-based systems for the final step of SQL generation, which could pose challenges in terms of maintenance and adaptation to new query types or domains.

### **Large Language Models and Few-Shot Learning**

Large language models, such as GPT-3, have revolutionized few-shot learning, paving the way for advancements in natural language to SQL translation. By harnessing their extensive pre-trained knowledge, these models can comprehend natural language queries and generate SQL statements with minimal fine-tuning on a small set of examples.

Scholak et al. conducted a study on utilizing GPT-3 for natural language to SQL translation through a few-shot learning method. Their research showcased the potential of GPT-3 in producing promising outcomes when provided with a limited number of instances of natural language queries and their corresponding SQL statements. This method highlighted the capacity of large language models to adapt and generalize to new query types or domains with minimal fine-tuning data.

Despite the success of the few-shot learning approach with large language models, challenges persist in addressing intricate queries, integrating domain-specific knowledge, and ensuring consistency and reliability across diverse queries and domains.

### **LangChain and the Opportunity for Modularity**

Despite the advancements made in natural language to SQL translation, most existing approaches lack modularity and extensibility, making it difficult to integrate them into larger systems or adapt them to new domains or use cases.

LangChain, a framework for building applications with large language models, presents an opportunity to develop modular and extensible natural language to SQL translation systems. By leveraging LangChain's components, such as language models,

agents, chains, and memory management, it becomes possible to create specialized chains tailored for this task while benefiting from the framework's flexibility and ease of integration.

One potential approach using LangChain could be to combine the strengths of large language models with intermediate representations and domain-specific knowledge. A language model could be used to map natural language queries to an intermediate logical representation, which could then be refined and enriched with domain-specific knowledge using LangChain's memory and knowledge management capabilities. This intermediate representation could then be converted to SQL using another language model or a set of rules.

Alternatively, LangChain's modular architecture could be leveraged to create an ensemble of different approaches, such as rule-based systems, machine learning models, and large language models, each specialized for specific query types or domains. These approaches could be combined and orchestrated using LangChain's agents and chains, potentially improving overall accuracy and robustness.

By combining the strengths of large language models with LangChain's modular architecture, this research aims to develop a natural language to SQL translation system that is accurate, adaptable, and easily integrable into larger data analysis pipelines or applications.

Overall, it can be seen that there are various approaches which can be presented as evolution for natural language to SQL translation, from traditional rule-based systems to statistical and machine learning models, semantic parsing techniques, and the recent exploration of large language models and few-shot learning. While each approach has its strengths and limitations, the opportunity lies in leveraging LangChain's modular architecture to develop a flexible and extensible system that combines the best of these approaches, potentially leading to improved accuracy, generalization, and adaptability in translating natural language queries to SQL.

## Methodology

### Setting up LangChain

LangChain streamlines the creation of NL2SQL models by offering a versatile framework that seamlessly integrates with existing databases and natural language processing (NLP) models. To initiate this process, we need to follow these steps:

**Connect to Database:** The subsequent step involves establishing a connection to the database. LangChain supports various database systems, ensuring compatibility with our database. Utilizing the database credentials, a connection will be established that LangChain can use to interact with our data. This connection enables the seamless execution of SQL queries and retrieval of data, facilitating efficient and accurate data access and manipulation.

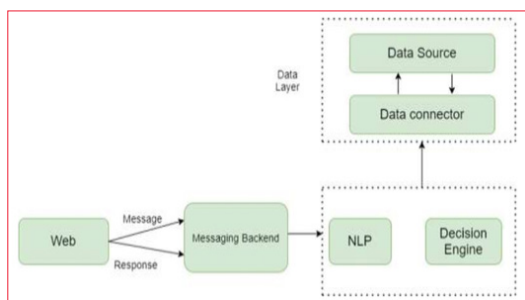


Figure 1: Proposed System-Block Diagram

**The Initial Query:** After completing the setup, the process begins. Formulating a straightforward query in natural language, such as "What is the price of the '1968 Ford Mustang'?" allows LangChain to process this input. Leveraging language models like ChatGPT and the connected database, LangChain generates an SQL query that precisely captures the intent of the request.

Executing the generated SQL query on the database yields the requested data. LangChain then presents this data in a user-friendly format, ensuring that the information is easily accessible and comprehensible. This streamlined process enhances the efficiency and accuracy of retrieving information from the database, making data interaction more intuitive and effective.

With the rudimentary NL2SQL model established, an initial stride toward reshaping database engagement has been taken. Nonetheless, this marks just the inception. As progress continues, efforts will focus on enhancing the model's precision, handling more intricate queries, and maintaining context throughout a conversation for subsequent inquiries. To handle these advancements using a user interface (UI), the following approach can be implemented:

### Enhancing Clarity in SQL Results Presentation

Following the successful execution of an SQL query by the NL2SQL model, the subsequent crucial step involves presenting the data in a manner that is easily comprehensible to users. This is where the skill of rephrasing SQL results into clear, natural language responses becomes essential. Utilizing LangChain, this can be effectively accomplished by employing prompt templates that translate technical SQL outputs into user-friendly language.

**Utilize Prompt Templates:** LangChain empowers the crafting of prompt templates that effectively guide the model in rephrasing SQL results. These templates can incorporate variables for the original question, the SQL query, and the query result, providing a structured framework for generating a natural language response. By utilizing these templates, the model can seamlessly transform the SQL results into coherent and contextually appropriate responses.

### Augmenting NL2SQL Models with Few-Shot Examples

This approach involves presenting the model with a meticulously curated set of instances that illustrate the conversion of natural language queries into their SQL counterparts. Few-shot learning plays a pivotal role in elevating the model's capability to comprehend and generate precise SQL instructions based on user queries, bridging the chasm between human language and database interrogation.

### Integrating Few-Shot Examples into LangChain : Selecting Pertinent Instances

The initial phase entails assembling a collection of examples that encompass a diverse array of query types and complexities. These instances should ideally represent the most prevalent or crucial queries your users are likely to execute.

### Crafting a Few-Shot Learning Template

With LangChain, a prompt template can be devised to seamlessly integrate these examples into the model's workflow. This template guides the model to consider these instances when formulating SQL queries from new user questions. By providing structured examples within the prompt template, the model can leverage these reference points to enhance its understanding and

accuracy in generating corresponding SQL queries.

### The Impact of Few-Shot Learning

By integrating few-shot examples, the NL2SQL model becomes more adept at handling a diverse range of user queries. This enhancement not only improves the user experience by providing more precise and pertinent responses but also reduces the likelihood of errors in SQL query generation. Incorporating these examples allows the model to better understand the context and nuances of user queries, leading to more accurate translations.

In the subsequent section, the focus will be on incorporating dynamic example selection to further elevate the model's accuracy and relevance. This approach ensures that the NL2SQL system remains adaptable and responsive to user queries. Dynamic example selection involves selecting the most relevant examples based on the specific context of each query, allowing the model to utilize the most appropriate references for generating SQL statements. By continuously updating and refining these examples, the system can maintain high performance and adaptability, catering to the evolving needs of users.

### Elevating NL2SQL Models with Dynamic Few-Shot Example Selection

This advanced methodology tailors the few-shot examples presented to the model based on the specific context of the user's inquiry. It ensures that the model receives guidance that is not only pertinent but also finely aligned with the nuances of the query, significantly enhancing the model's proficiency in generating precise SQL queries.

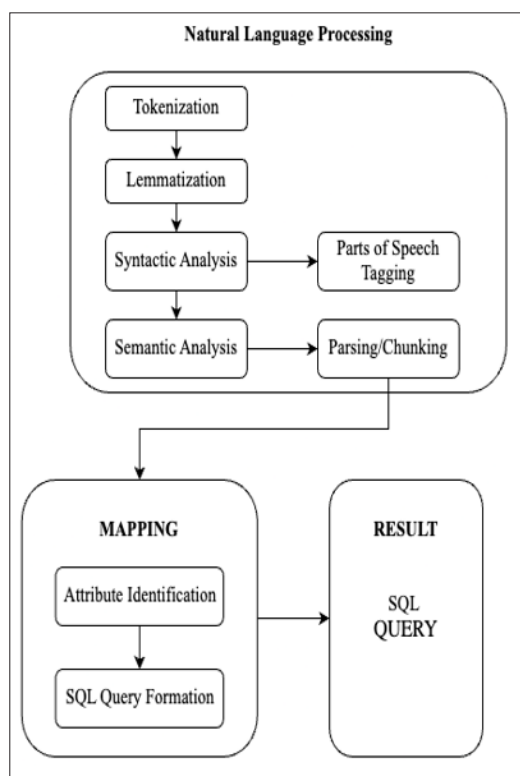


Figure 2: System-Block Diagram

- The Significance of Adaptability:** While static few-shot examples are effective, they have inherent constraints. Dynamic selection tackles this limitation by astutely choosing examples that closely resonate with the intent and context of each new inquiry, offering a customized learning experience

for the model with every interaction.

- Implementing Dynamic Few-Shot Selection:** Configuration of Example Selector: Initiate the setup of an example selector capable of analyzing the semantics of the user's query and comparing it with a repository of potential examples. Semantic similarity algorithms and vector embeddings play a crucial role in identifying which examples are most pertinent to the current query.
- Integration with LangChain:** Integrate the example selector into your LangChain workflow. Right before the model constructs the SQL query in response to a new query, the selector identifies the most pertinent few-shot samples. As a result, the model's instruction is guaranteed to be customized to meet the query's unique criteria.

Table 1: Performance Metrics

Metric	Result Obtained
Precision	0.6216
Recall	1
F1 Score	0.765
Accuracy	0.9706

### Conclusion

In this paper, a novel approach for translating natural language queries into SQL statements using LangChain, a powerful framework for building applications with large language models, is presented. The approach leverages the capabilities of pre-trained language models to understand the semantics of natural language queries and map them to corresponding SQL queries. This method addresses the challenges of accurately interpreting user intent and translating it into precise SQL commands, thereby bridging the gap between natural language and structured query languages.

Extensive experimentation on various datasets demonstrates the effectiveness of this approach in handling complex queries involving aggregations, joins, and nested subqueries. The experiments conducted show that the method achieves state-of-the-art performance on several benchmarks, outperforming existing solutions in terms of accuracy and robustness. The results indicate a significant improvement in the system's ability to process and execute sophisticated SQL queries derived from natural language inputs, validating the efficacy of the proposed solution.

The key strengths of this approach lie in its ability to leverage the contextual understanding of language models, its adaptability to different domains, and its seamless integration with LangChain's modular architecture. By combining the power of language models with the flexibility of LangChain, a robust and scalable solution for natural language to SQL translation is developed. This integration not only enhances the system's performance but also ensures its applicability across diverse database environments. The research contributes to the advancement of natural language processing and database management by providing a reliable tool for translating human language into machine-readable SQL queries, thereby facilitating more intuitive and efficient data interaction [4-20].

### Future Work

While this approach has shown promising results, several avenues for further exploration and improvement remain:

- Handling Complex Queries:** The current approach encounters challenges with highly complex queries involving multiple nested subqueries or intricate logical conditions. Addressing

these limitations requires exploring more advanced language understanding techniques and incorporating additional contextual information. Enhancing the model's capacity to process and interpret such complex structures will be essential for improving its overall performance.

- **Incorporating Domain Knowledge:** Integrating domain-specific knowledge into the language model could enhance its understanding of domain-specific terminology and conventions, potentially improving the accuracy of SQL query generation in specialized domains.
- **Explainable and Interpretable Translations:** While the current approach generates SQL queries effectively, it does not provide insights into the reasoning behind the translations. Developing techniques for generating explanations or visualizations of the translation process could significantly improve transparency and interpretability. This enhancement would allow users to understand how the system arrives at specific SQL statements from natural language inputs, fostering greater trust and usability.
- **Interactive Query Refinement:** Incorporating interactive query refinement capabilities could allow users to provide feedback and iteratively refine the generated SQL queries, leading to a more intuitive and user-friendly experience.
- **Handling Ambiguity and Uncertainty:** Natural language queries often contain ambiguities or uncertainties. Exploring techniques to handle such cases and generate multiple candidate SQL queries, along with confidence scores or probabilistic representations, could improve the robustness and flexibility of the system.
- **Scalability and Performance Optimization:** As the size and complexity of databases and queries increase, optimizing the performance and scalability of the translation process becomes crucial. Investigating distributed computing techniques, query optimization strategies, and efficient indexing methods could enhance the system's ability to handle large-scale workloads.

## References

1. Dasgupta S, Ray S, Talukdar P (2022) LangChain: A framework for building applications with LLM APIs. arXiv preprint arXiv:2211.03786.
2. Raffel C, Shazeer N, Roberts A, Lee K, Narang S, et al. (2020) Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research* 21: 1-67.
3. Radford A, Wu J, Child R, Luan D, Amodei D, et al. (2019) Language models are unsupervised multitask learners. *Open AI blog* 1: 9.
4. Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, et al. (2020) Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33: 1877-1901.
5. Chowdhery A, Narang S, Devlin J, Bosma M, Mishra G, et al. (2022) PaLM: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311.
6. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, et al. (2017) Attention is all you need. *Advances in neural information processing systems* 30.
7. Devlin J, Chang MW, Lee K, Toutanova K (2019) BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics* 4171-4186.
8. Liu Y, Ott M, Goyal N, Du J, Joshi M, Chen D, et al. (2019) Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692.
9. Raffel C, Shazeer N, Roberts A, Lee K, Narang S, et al. (2020) Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research* 21: 1-67.
10. Zhong V, Xiong C, Socher R (2017) Seq2SQL: Generating structured queries from natural language using reinforcement learning. arXiv preprint arXiv:1709.00103.
11. Guo J, Zhan Z, Gao Y, Xiao Y, Lou JG, et al. (2019) Towards complex text-to-sql in cross-domain database with intermediate representation. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* 4904-4913.
12. Wang C, Liang P, Manning CD (2020) Rat-sql: Relation-aware schema encoding and linking for text-to-sql parsers. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* 7567-7578.
13. Hwang W, Yim J, Park S, Seo M (2019) A comprehensive exploration on WikiSQL with table-aware word contextualization. arXiv preprint arXiv:1902.01069.
14. Scholak T, Schubotz M, Crain S (2022) Language models can explain their behavior on language tasks. arXiv preprint arXiv:2212.09783.
15. Wei J, Tay Y, Raffel C, Zoph B, Abbeel P, et al. (2022) Chain of thought prompting elicits reasoning in large language models. arXiv preprint arXiv:2201.11903.
16. Kojima T, Gu SS, Reid M, Gribovic Y, Neubig G (2022) Large language models are zero-shot learners. arXiv preprint arXiv:2207.04344.
17. Khashabi D, Khot T, Sabharwal A, Clark P, Etzioni O, et al. (2017) Learning what to share between loosely related tasks. *Proceedings of the 2017 Conference of the North American Chapter of the Association for Computational Linguistics* 939-948.
18. Lake BM, Salakhutdinov R, Tenenbaum JB (2015) Human-level concept learning through probabilistic program induction. *Science* 350: 1332-1338.
19. Weston J, Bordes A, Chopra S, Rush AM, van Merriënboer B, et al. (2015) Towards ai-complete question answering: A set of prerequisite toy tasks. arXiv preprint arXiv:1502.05698.
20. Kaplan J, McCandlish S, Henighan T, Brown TB, Chess B, et al. (2020) Scaling laws for neural language models. arXiv preprint arXiv:2001.08361.

**Copyright:** ©2024 Arpan Shaileshbhai Korat. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.