

The Future of Database Systems: Innovations and Challenges in Natural Language Interfaces

Manikkaarachchi R. N
Griffith University , QLD,
Australia

Abstract:- Natural Language Interfaces (NLIs) have significantly improved accessibility to database systems by allowing users to interact using natural language queries rather than complex query languages. This paper examines recent advancements in NLP and machine learning that enhance NLI functionality for database systems, discusses current methodologies and technologies, addresses the major challenges these systems face, and proposes future research directions. NLIs have broad applications, including business intelligence and customer service, where simplifying database access can streamline operations and support data-driven decision-making.

Keywords:- Natural Language Interface (NLI), Deep Learning, Contextual Understanding, Conversational AI, Structured Query Language (SQL), ML (Machine Learning).

I. INTRODUCTION

With the rapid expansion of data generation by organizations, there has been an increased demand for sophisticated database systems for efficient storage, retrieval, and analysis. Traditionally, these databases required expertise in query languages, like SQL, to interact with the data effectively. However, such technical expertise is not always accessible to all users, creating a barrier to data access and analysis for non-technical stakeholders.

Natural Language Interfaces (NLIs) to databases have been developed as a solution, enabling users to interact with databases using everyday language. NLIs convert natural language queries into structured database queries, allowing users to obtain data insights without needing SQL proficiency. The relevance of NLIs spans multiple sectors, including healthcare, business intelligence, and customer service, where timely data retrieval is crucial. Despite promising developments, NLIs still face challenges in handling complex queries, maintaining context over multiple queries, and ensuring accurate query interpretation.

This paper explores the advancements in NLP and machine learning that have facilitated NLI development, reviews current methodologies for query generation and error handling, and identifies the challenges that remain. Additionally, it outlines potential future directions to improve the usability, accuracy, and scalability of NLIs for database system.

II. LITERATURE REVIEW

A. Early Developments and Rule-Based Systems

The first generation of Natural Language Interfaces (NLIs) emerged in the late 1970s and early 1980s, when rule-based and template-driven systems were widely used. These systems relied on predefined syntactic templates and lexical mappings that matched common query patterns. A notable example is the LUNAR system developed by Woods in 1973, designed to query a database of lunar rock samples using natural language commands. LUNAR translated English sentences into a structured database query language by matching sentence patterns to grammar rules. Despite its limitations, LUNAR demonstrated the potential of NLIs for non-technical users, but it was highly domain-specific and struggled to handle unforeseen language constructions.

Another example from this era is the LADDER system, which utilized a grammar-based approach to interpret and translate user queries into structured database queries. LADDER defined syntactic rules to parse English sentences, enabling users to retrieve geographic information without technical training. While this rule-based system allowed for some flexibility, it was limited by a rigid rule set that struggled with ambiguous or novel phrasing (Woods, 1973; Androutsopoulos et al., 1995).

B. Statistical and Template-Based NLIs

In the 1990s and early 2000s, template-based and statistical models began to address some limitations of rule-based systems by focusing on probabilistic methods. Template-based approaches predefined mappings from phrases to database queries, allowing the system to handle common patterns more reliably. For example, TEAM by Thompson et al. (1985) used templates to process natural language questions by matching them to specific database commands. Though an improvement, template-based systems lacked flexibility for dynamic and evolving query types (Thompson et al., 1985).

Statistical approaches as Hidden Markov Models (HMMs) and Conditional Random Fields (CRFs) introduced probabilistic models for language processing, marking a shift towards more adaptable NLIs. These models were applied in systems like ELIZA (Weizenbaum, 1966) and SHRDLU (Winograd, 1972), which, while offering limited conversational capabilities, could interact with databases through pattern matching. However, these systems were not designed to handle complex data retrieval tasks. Statistical

NLIs contributed to later machine learning approaches by emphasizing probability in language parsing but struggled with scalability for larger databases (Jurafsky & Martin, 2009).

C. Emergence of Semantic Parsing and Machine Learning Approaches

The early 2000s saw a shift toward semantic parsing, which aimed to understand the meaning behind words rather than just their syntactic placement. Semantic Parsing systems convert natural language sentences into formal representations by identifying entities and their relationships, making it easier to map sentences to database queries. One of the foundational works was by Zelle and Mooney (1996), who developed GeoQuery, a system that could interpret natural language questions about U.S. geography. GeoQuery used inductive logic programming to generalize from examples, demonstrating that semantic parsing could support more sophisticated NLIs (Zelle & Mooney, 1996).

Building on these advancements, Wang et al. (2015) introduced the STAGG model, which leveraged semantic parsing with alignment to generate SQL queries from natural language input. By learning from example sentences, STAGG was able to adapt to user queries more dynamically, representing a significant improvement over earlier approaches. This model marked a notable step forward in the development of natural language interfaces for database querying.

D. Advancements in Deep Learning Approaches for Natural Language Interfaces

The advent of deep learning, particularly sequence-to-sequence (Seq2Seq) models, has led to significant advancements in Natural Language Interfaces (NLIs) by enhancing the ability to interpret complex queries. Initially designed for machine translation (Sutskever et al., 2014), Seq2Seq models are well-suited for translating natural language into SQL queries due to their capacity to process input sequences and generate variable-length output sequences.

Seq2SQL (Zhong et al., 2017) introduced reinforcement learning to further improve the accuracy of translating natural language queries into SQL. Seq2SQL was trained on large datasets, learning how to generate SQL code and refining its predictions based on reward signals tied to accuracy. This model achieved significant success on the WikiSQL benchmark, a large-scale dataset that pairs natural language questions with SQL queries (Zhong et al., 2017).

The transformer model, introduced by Vaswani et al. (2017), advanced NLIs for databases by capturing contextual relationships across sentences and phrases more effectively than previous models. Transformer-based models, such as BERT (Bidirectional Encoder Representations from Transformers) and GPT (Generative Pre-trained Transformer), have demonstrated remarkable success in question answering and text generation, making them highly applicable for NLIs. For example, IRNet (Guo et al., 2019) utilized transformers to generate SQL queries from natural

language with high accuracy by understanding complex relationships between table attributes and query components (Guo et al., 2019).

E. Conversational NLIs and Multi-Turn Dialogues

Modern NLIs part multi-turn dialogue, where users interact with the system through a series of exchanges. This feature enables users to refine their queries iteratively, reducing ambiguity and improving result accuracy. Microsoft's Power BI Q&A and IBM Watson are notable examples of conversational AI systems that support complex database queries through dialogue. Research by Li and Jagadish (2014) on conversational interfaces highlighted the potential of multi-turn dialogues to resolve ambiguities and enhance user satisfaction (Li & Jagadish, 2014).

F. Benchmarking and Datasets

Evaluation of metrics and benchmarking datasets essential for assessing NLI performance. Popular benchmarks like WikiSQL (Zhong et al., 2017) and Spider (Yu et al., 2018) provide standardized datasets for evaluating NLI systems on SQL query generation. These datasets allow researchers to compare models across standardized tasks, fostering improvements in NLI accuracy and robustness (Yu et al., 2018).

III. METHODOLOGIES IN NLI FOR DATABASES

Natural Language Interfaces (NLIs) to databases rely on multiple methodologies to accurately interpret and translate natural language into structured database queries. These methodologies include lexical and semantic parsing, query generation, and error handling, all of which contribute to effective and reliable database querying.

A. Lexical and Semantic Parsing

Lexical parsing, the process of tokenizing a user query into individual words or phrases, is the foundation for understanding the components of a query. Semantic parsing expands on this by interpreting the relationships and meanings behind these tokens, transforming them into a structured format like SQL. Early approaches in semantic parsing involved grammar-based systems and rule-based templates (Woods, 1973; Androutsopoulos et al., 1995), which, while functional, struggled to generalize across domains. In recent years, advances in transformer models and attention mechanisms have significantly enhanced the ability of NLIs to interpret context and relationships within queries, improving overall parsing accuracy (Vaswani et al., 2017).

Techniques such as dependency parsing and semantic role labelling play crucial roles in these processes. Dependency parsing identifies grammatical relationships between words, clarifying their roles and dependencies within a sentence. This technique has been effective in NLI models, as it aids in understanding how various tokens connect to form meaningful database queries (Li & Jagadish, 2014). In contrast, semantic role labeling assigns specific roles to parts of a sentence, such as identifying the subject, action, and object, thereby helping the system generate

accurate and contextually appropriate queries (Jurafsky & Martin, 2009).

B. Query Generation

Once a system has parsed the natural language input, it must convert this interpretation into a structured query, typically SQL. Template-based query generation was among the earliest methods used, where predefined templates mapped specific phrases to SQL commands (Thompson et al., 1985). However, template-based systems lack flexibility, particularly in handling complex queries. Rule-based systems offered slightly more adaptability but still fell short with novel language patterns, leading researchers to explore machine learning models for query generation.

The introduction of neural machine translation (NMT) models, such as sequence-to-sequence (Seq2Seq) architectures, marked a significant advancement. Seq2Seq models can learn mappings from natural language to SQL, a concept demonstrated by Seq2SQL, which combined sequence generation with reinforcement learning to refine SQL predictions based on accuracy (Zhong et al., 2017). Transformer-based architectures, such as BERT and GPT, have further enhanced query generation capabilities by capturing the nuances in user input and translating them into accurate database commands (Guo et al., 2019). These models benefit from large datasets, such as WikiSQL and Spider, which provide valuable training examples that improve model generalizability across domains (Yu et al., 2018).

C. Error Handling and Ambiguity Resolution

Handling ambiguous or erroneous queries is a crucial part of an effective NLI. Systems must identify when a user's input lacks clarity or completeness and guide them toward refinement. Techniques such as query rewriting and interactive dialogue systems have been implemented to address this issue. Query rewriting suggests alternative phrasings or corrections to help users clarify their intent, while dialogue-based systems engage users in multi-turn conversations to resolve ambiguities, especially useful for complex or incomplete queries (Li & Jagadish, 2014). Additionally, confidence scoring helps the system gauge the reliability of its query interpretation, prompting user feedback if confidence is low. This iterative approach enables NLIs to handle ambiguity more effectively and improve accuracy over time through user interaction (Yu et al., 2018).

IV. CURRENT APPROACHES AND TECHNOLOGIES

Recent advancements in Artificial Intelligence (AI) and deep learning have driven the development of more sophisticated Natural Language Interfaces (NLIs) for databases. Modern systems leverage powerful models like transformers, BERT, and GPT to improve query interpretation and contextual understanding, while conversational interfaces further enhance database accessibility. This section outlines these cutting-edge approaches and technologies.

A. AI and Deep Learning in NLI

AI and deep learning models have transformed the capabilities of NLIs, especially with the introduction of transformer architectures like BERT (Bidirectional Encoder Representations from Transformers) and GPT (Generative Pre-trained Transformer). These models excel at interpreting complex, nuanced language due to their ability to capture context and relationships in text. Transformers leverage attention mechanisms to model long-range dependencies in language, which is crucial for understanding the structure and meaning of database queries (Vaswani et al., 2017).

BERT has been widely used for language understanding tasks due to its bidirectional approach, which allows it to consider context from both the left and right sides of a target word. BERT's pre-trained contextual embeddings have proven effective for parsing database queries by capturing intricate language details, improving the accuracy of query generation and interpretation (Devlin et al., 2019). GPT, a generative model, has also shown promise in NLI for its capacity to generate structured queries based on context-rich prompts. The model's autoregressive architecture generates text by predicting each word sequentially, making it useful for handling complex and open-ended user queries (Brown et al., 2020).

Seq2Seq models and attention-based neural networks are also frequently used in text-to-SQL tasks, transforming natural language questions into structured SQL queries. For example, Seq2SQL uses reinforcement learning to refine query predictions and improve response accuracy (Zhong et al., 2017). Recent research also integrates transfer learning and fine-tuning on domain-specific data, making it easier to adapt these deep learning models to specialized applications.

B. Contextual Understanding

Contextual understanding is essential for interpreting queries accurately, especially in systems that allow for sequential queries or conversations. Modern NLIs use contextual embeddings to capture and retain the information from previous queries, enabling the system to infer relationships and dependencies between queries over time. This feature is particularly beneficial in applications where users build on previous questions, as it allows the system to maintain coherence across sessions without requiring the user to restate prior information explicitly.

For instance, transformer-based architectures retain and apply context within conversational threads, allowing the NLI to interpret follow-up queries with improved accuracy. Contextual models like BERT can handle follow-up questions by considering the information from prior interactions, leading to more accurate responses. Additionally, memory-augmented models are increasingly being used to track session-based data, enhancing the ability of NLIs to maintain continuity across multiple queries (Su et al., 2020). By providing contextual awareness, these systems better mimic natural, human-like dialogue patterns and support more complex database interactions.

C. Conversational Interfaces

Conversational interfaces bring a more interactive and accessible approach to NLIs, allowing users to query databases in a way that resembles a natural conversation. These interfaces enable iterative querying, allowing users to refine questions, request clarifications, or adjust parameters mid-query. Through conversational AI, users can pose complex questions more intuitively and receive responses in a manner that resembles a dialogue, enhancing user experience and reducing the complexity of query construction.

Prominent tools like Microsoft Power BI Q&A, IBM Watson, and Google Big Query ML have made substantial progress in providing conversational access to databases. Microsoft Power BI Q&A leverages natural language processing (NLP) to allow users to ask questions in plain language, and it provides visual responses based on data insights. It interprets user intent and automatically generates visualizations, making it highly accessible for non-technical users (Microsoft, n.d.). IBM Watson incorporates conversational AI and NLP, enabling businesses to build NLIs that support context-driven interactions and clarify ambiguous queries through follow-up questions (IBM, n.d.). In contrast, Google Big Query ML integrates machine learning capabilities directly into its database, allowing for SQL-based queries augmented by NLP functions, facilitating both structured and unstructured data analysis in conversational formats (Google, n.d.).

These conversational interfaces simplify access to complex data and provide a more interactive experience. They are particularly useful in business intelligence, where real-time, natural language querying supports data-driven decision-making without the need for SQL expertise.

V. CHALLENGES AND LIMITATIONS

Despite significant progress, Natural Language Interfaces (NLIs) for databases still face notable challenges and limitations that impact their effectiveness and adoption. These issues range from language complexity to data heterogeneity, domain specificity, and user interaction constraints.

A. Language Complexity and Ambiguity

A key challenge for Natural Language Interfaces (NLIs) lies in managing the complexity and inherent ambiguity of natural language. Users often phrase queries in varied and complex ways, making it difficult for an NLI to interpret the intended meaning accurately. Ambiguity can arise from vague or polysemous words, where a word has multiple meanings (e.g., "bank" as a financial institution or riverbank), or from unclear sentence structures. While advances in deep learning and context-aware models like transformers have improved the ability to capture nuanced language meanings, accurately resolving ambiguity in all cases remains challenging. This limitation affects both lexical and semantic parsing, often leading to incorrect query translations and a need for more robust disambiguation techniques.

B. Domain Adaptability and Specificity

Most NLI models are trained on specific datasets and domains, making them highly optimized for types of queries but less effective when applied to different contexts. For instance, a system trained on medical databases may struggle with financial data or geographical queries. Domain specificity limits scalability because each new application often requires extensive retraining and fine-tuning on domain-specific data. This is particularly challenging for NLIs in enterprise applications, where databases may contain highly specialized or proprietary terminology. Research on transfer learning and domain adaptation has sought to address this issue, but generalized, cross-domain NLI capabilities are still largely under development.

C. Handling Complex Queries and Multi-Table Joins

NLIs can struggle with complex queries, especially those that involve multiple tables, nested clauses, or complex joins. Translating these types of queries into SQL is a non-trivial task that requires the system to understand intricate relationships between tables and interpret multi-layered conditions. Current systems, including those leveraging deep learning, often face accuracy drops when handling complex query structures. Although recent models, such as transformer-based architectures, have improved performance on complex queries, the challenge of maintaining high accuracy across all levels of query complexity persists. As a result, complex database queries still require manual intervention or refinement.

D. Lack of Explainability

Explainability remains a significant challenge for NLIs, particularly those using deep learning methods. Models such as Seq2Seq and transformers produce high-quality translations of natural language into SQL, but they do so in a black-box manner. This opacity makes it difficult for users and developers to understand why a particular SQL query was generated, especially in cases where the result is incorrect or unexpected. This lack of transparency limits trust and reliability in NLIs, particularly in high-stakes fields such as healthcare or finance, where understanding the basis of decisions is crucial. Efforts to build explainable AI, such as integrating interpretability layers or visualizations, are still nascent within the NLI domain.

E. Dependency on High-Quality Training Data

NLI systems rely heavily on high-quality training data to perform effectively. However, labeled datasets that map natural language questions to database queries are limited, especially for complex queries or specialized domains. The WikiSQL and Spider datasets are notable for SQL generation, but they are insufficient for covering the full range of possible database interactions and domain-specific nuances. This lack of comprehensive, labeled data hampers the ability of NLIs to generalize across domains or adapt to novel queries. Moreover, creating and annotating such data is time-consuming and costly, posing a barrier to broader adoption and improvement of NLIs.

F. User Interaction and Feedback Mechanisms

While NLI systems strive to make database querying accessible, they often lack effective mechanisms for managing user interactions and feedback. User queries may be incomplete or incorrect, and systems without feedback loops struggle to improve from such interactions. Dialog-based NLIs, which engage users in iterative query refinement, have been proposed to address this, but implementing robust multi-turn dialogue capabilities is challenging and computationally intensive. Many NLIs also lack mechanisms to guide users through query refinement or explain why certain query interpretations failed, resulting in a less user-friendly experience.

G. Real-Time Performance Constraints

For large databases, real-time query processing can be a bottleneck for NLIs. The translation of natural language into SQL often involves multiple computational steps, including parsing, semantic analysis, and query generation, all of which must be completed quickly to deliver an interactive experience. Systems that rely heavily on deep learning may face latency issues, particularly when processing complex queries or interacting with massive datasets. Performance constraints can deter adoption, as users may expect immediate responses akin to conversational AI, while back-end systems may struggle to keep up with real-time demands.

VI. FUTURE DIRECTIONS

Addressing the challenges faced by current Natural Language Interfaces (NLIs) to databases offers promising directions for future development. Future advancements are likely to focus on improving interpretability, domain adaptability, handling complex queries, integrating richer contextual understanding, and enhancing user interactions.

A. Enhanced Explainability and Interpretability

As deep learning models become more complex, enhancing explainability remains a priority. Future NLIs could integrate explainable AI (XAI) techniques that offer insights into how queries are interpreted and transformed into SQL. For example, attention visualization and feature attribution methods could help users understand which parts of a question influenced the generated query. Explainable NLI models will be especially valuable in sensitive fields like healthcare and finance, where accountability and transparency are critical. Research into making black-box models more interpretable is ongoing, with methods such as layer-wise relevance propagation (LRP) and shapley additive explanations (SHAP) showing promise for making model decisions more transparent (Samek et al., 2017).

B. Domain Adaptability and Transfer Learning

To overcome limitations in domain specificity, future NLIs could leverage transfer learning and meta-learning techniques, which allow models to adapt to new domains with minimal retraining. By using large-scale, domain-agnostic models and fine-tuning them on domain-specific data, NLIs could achieve greater flexibility across diverse fields. This approach would enable NLIs to generalize effectively, reducing the need for extensive re-training on

every new dataset. Furthermore, few-shot learning techniques could allow NLIs to handle rare or new queries by drawing on patterns from limited domain-specific data, thus enhancing adaptability and reducing the demand for domain-specific datasets (Hospedales et al., 2021).

C. Improved Handling of Complex and Multi-Table Queries

The ability to handle complex, multi-table joins and nested queries remains a challenging area. Future NLIs could benefit from multi-hop reasoning models that process queries in steps, building intermediate representations that can accurately handle multi-table joins and conditional statements. Graph-based representations of databases might also be used to improve the system's understanding of table relationships and hierarchies, making it easier to generate complex SQL queries. Research into knowledge graph embeddings and graph neural networks (GNNs) is continually showing promise in enhancing relational understanding, which could translate to improved performance in complex database querying (Wu et al., 2020).

D. Advanced Contextual Understanding with Memory-Augmented Models

Contextual understanding in NLIs can be enhanced by using memory-augmented neural networks that retain and leverage information across interactions. Such models can be particularly useful in conversational NLIs, where each query may build on previous ones. Memory networks, along with transformer-based architectures that support session-based data retention, could help NLIs maintain coherence across extended interactions. Incorporating session-based memory will allow future systems to support complex workflows, enabling users to iteratively refine queries or return to previous contexts without loss of information, which could improve usability and result accuracy.

E. Human-in-the-Loop for Improved Interaction and Feedback

A human-in-the-loop approach can further improve NLI systems by incorporating user feedback into model training. This approach allows NLIs to learn continuously from user corrections, adjusting interpretations and improving accuracy over time. Interactive feedback mechanisms, such as prompts for clarification when confidence is low or suggestions for alternative query formulations, could also enhance user experience. In addition, dialogue-based systems that offer multi-turn conversations to refine ambiguous queries are likely to become more prevalent, as they allow the NLI to handle errors more naturally and provide real-time clarification, improving the overall usability of NLIs (Kumar et al., 2020).

F. Integration with Conversational and Augmented Reality Interfaces

The integration of NLIs with conversational interfaces and augmented reality (AR) could revolutionize how users interact with complex data. Conversational agents, particularly those leveraging technologies like Microsoft Power BI Q&A and IBM Watson, are increasingly providing voice-based and real-time responses, making it easier for non-technical users to query and explore databases.

Augmented reality interfaces could take this further by offering hands-free, interactive experiences where users can verbally query databases and receive visualized data overlays in real-time. This approach has potential applications in fields such as healthcare, where clinicians could retrieve and interpret data directly from patient charts via AR glasses during treatment (McIntire et al., 2018).

G. Cross-Language Capabilities for Multilingual Querying

As databases become global, NLI with cross-language querying capabilities will be essential. Future systems could integrate multilingual NLP models that support natural language queries in multiple languages, enabling cross-cultural accessibility to databases. Models like mBERT (Multilingual BERT) and XLM-R (XLM-RoBERTa) already support cross-language understanding and could be extended for NLI applications, enabling users worldwide to query databases in their native language. Multilingual capabilities will require advancements in cross-lingual embeddings and translation techniques, particularly for specialized vocabulary and domain-specific terms.

H. Real-Time Performance Optimization

Finally, as NLIs interact with larger databases, optimizing real-time performance will be critical. Efficient model architectures and hardware optimizations, such as those leveraging edge computing and model distillation, will be necessary to ensure that NLIs can provide immediate responses, even with computationally intensive tasks. Techniques like quantization and pruning are already being applied to reduce model sizes without sacrificing performance, and these methods could be adapted to improve NLI response times for large-scale applications (Han et al., 2016).

VII. CONCLUSION

The field of Natural Language Interfaces (NLIs) for databases has seen substantial advancements, due to developments in artificial intelligence, deep learning, and natural language processing. These systems have made data access more intuitive, allowing users to interact with databases in natural language rather than requiring technical query languages like SQL. Current approaches, leveraging transformer models such as BERT and GPT, enhance query interpretation by capturing contextual meanings and understanding complex query structures. Additionally, conversational interfaces from platforms like Microsoft Power BI Q&A, IBM Watson, and Google Big Query ML have made NLIs more accessible, promoting a more interactive data querying experience.

Despite these advances, several challenges remain. NLIs must tackle issues related to context retention, interpretability, and handling of complex, multi-table queries, especially in specialized or highly regulated domains. Achieving domain adaptability, enhancing contextual memory, and improving cross-language support are crucial for extending the applicability of NLIs across industries. Moreover, creating explainable and transparent models will be essential for building user trust, particularly in sectors

where data-driven decisions directly impact human lives, such as healthcare and finance.

Future directions in NLI research present exciting possibilities. Enhanced explainability, memory-augmented models for contextual continuity, human-in-the-loop interaction, and multilingual capabilities are set to transform NLIs into more versatile and user-friendly tools. Additionally, integrating NLIs with augmented reality and real-time optimization technologies will facilitate seamless, hands-free data interactions, further democratizing access to complex data.

In conclusion, while current NLIs provide a strong foundation, ongoing research and technological innovations will continue to push the boundaries of what these systems can achieve. By addressing existing limitations and exploring novel interaction paradigms, NLIs can evolve into indispensable tools for data-driven decision-making, empowering users across diverse fields to harness the full potential of their data.

REFERENCES

- [1]. Woods, W. A. (1973). Progress in Natural Language Understanding: An Application to Lunar Geology. *American Journal of Computational Linguistics*.
- [2]. Thompson, H. S., McDonald, D. D., & Wilensky, R. (1985). Getting to know the TEAM: A system for combining natural language understanding and planning.
- [3]. Androutsopoulos, I., Ritchie, G. D., & Thanisch, P. (1995). Natural language interfaces to databases - an introduction. *Journal of Language and Engineering*.
- [4]. Zelle, J. M., & Mooney, R. J. (1996). Learning to parse database queries using inductive logic programming. *Proceedings of AAAI-96*.
- [5]. Jurafsky, D., & Martin, J. H. (2009). *Speech and Language Processing*. Prentice Hall.
- [6]. Wang, C., Mazaitis, K., & Cohen, W. W. (2015). A Tensor-based Factorization Model for Domain-aware Semantic Parsing.
- [7]. Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in Neural Information Processing Systems (NIPS)*.
- [8]. Zhong, V., Xiong, C., & Socher, R. (2017). Seq2SQL: Generating structured queries from natural language using reinforcement learning. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- [9]. Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is All You Need. *Advances in Neural Information Processing Systems*.
- [10]. Li, F., & Jagadish, H. V. (2014). Constructing an Interactive Natural Language Interface for Relational Databases. *VLDB Endowment*.
- [11]. Guo, D., Sun, Y., Tang, D., et al. (2019). Towards complex text-to-SQL in cross-domain database with intermediate representation. *Proceedings of ACL 2019*.

- [12]. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. NAACL.
- [13]. Yu, T., Zhang, R., Zeng, Z., et al. (2018). Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task. EMNLP 2018.
- [14]. Samek, W., Wiegand, T., & Müller, K. R. (2017). Explainable artificial intelligence: Understanding, visualizing, and interpreting deep learning models. arXiv preprint arXiv:1708.08296.
- [15]. Brown, T. B., Mann, B., Ryder, N., et al. (2020). Language Models are Few-Shot Learners. NeurIPS.
- [16]. Wu, Z., Pan, S., Chen, F., et al. (2020). A comprehensive survey on graph neural networks. IEEE Transactions on Neural Networks and Learning Systems
- [17]. Kumar, A., Agarwal, S., & Luo, J. (2020). Dialogue-enabled interactive machine learning. ACM Transactions on Interactive Intelligent Systems.
- [18]. Su, C., Yan, X., Chen, T., et al. (2020). Contextual Sentence Classification in a Conversational AI Setting: A Performance Study on Transformer Models. IEEE Transactions on Affective Computing.
- [19]. Han, S., Mao, H., & Dally, W. J. (2016). Deep compression: Compressing deep neural networks with pruning, trained quantization, and huffman coding. International Conference on Learning Representations (ICLR).
- [20]. Hospedales, T. M., Antoniou, A., Micaelli, P., & Storkey, A. (2021). Meta-learning in neural networks: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [21]. McIntire, M., Miller, S., Li, J., & Smith, B. (2018). Augmented reality for medical applications: A systematic review. Journal of Medical Internet Research.
- [22]. Microsoft. (n.d.). Natural language queries in Power BI Q&A. Retrieved from <https://docs.microsoft.com>.
- [23]. IBM. (n.d.). IBM Watson Assistant. Retrieved from <https://www.ibm.com/watson/assistant>
- [24]. Google. (n.d.). BigQuery ML. Retrieved from <https://cloud.google.com/bigquery>