



Archives available at journals.mriindia.com

International Journal on Advanced Computer Engineering and Communication Technology

ISSN: 2278-5140

Volume 14 Issue 03s, 2025

Bridging Natural Language and Databases: Intelligent Conversational System for Automated SQL Query Generation Using Natural Language Processing

¹Avanti Gawande, ²Pranav Bhusari, ³Prabhu Ingole, ⁴Rushabh Wakekar

^{1,2,3,4} Department of Computer Science Engineering (Data Science), St. Vincent Pallotti College of Engineering and Technology, Nagpur, India

Email: ¹gawandeavanti09@gmail.com, ²bhusaripranav09@gmail.com, ³prabhuingole22@gmail.com, ⁴rushabhwakekar.22@stvincentngp.edu.in

Peer Review Information	Abstract
<p><i>Submission: 05 Nov 2025</i></p> <p><i>Revision: 25 Nov 2025</i></p> <p><i>Acceptance: 17 Dec 2025</i></p>	<p>Relational databases are the foundation of modern data management systems, widely utilised across industries for storing, organising, and retrieving information. However, interacting with these databases typically requires proficiency in Structured Query Language (SQL), which limits accessibility for non-technical users. Addressing this challenge, this paper presents a multimodal AI-powered SQL chatbot that allows users to query relational databases using natural language, voice commands, and visual schema inputs. The proposed system employs a fine-tuned T5 transformer model to convert natural language statements into executable SQL queries. It further integrates Optical Character Recognition (OCR) for schema extraction and multilingual translation modules to support diverse user inputs. Unlike conventional Text-to-SQL systems, this approach emphasises user-centric design, providing both textual and graphical result visualisation for improved interpretability. Experimental evaluation was conducted using benchmark datasets such as Spider and WikiSQL, alongside custom domain-specific datasets covering HR, Sales, and Academic data. The model achieved competitive accuracy and low latency, averaging 1.5–2 seconds per query, while ablation studies confirmed the impact of OCR and multilingual modules on performance. Furthermore, a user study involving 20 participants demonstrated high usability and satisfaction ratings, validating the chatbot’s potential to make database querying more natural and inclusive. This research contributes toward the democratisation of data access by merging transformer-based NLP, multimodal interaction, and user-focused design into an integrated, practical solution.</p>
<p>Keywords</p> <p><i>Natural Language Processing (NLP), SQL Chatbot, OCR, voice-driven.</i></p>	

Introduction

Relational databases serve as the backbone of nearly all modern information systems, providing structured storage, organisation, and retrieval of large-scale data. Despite their ubiquity, accessing information from these

systems generally requires specialised knowledge of Structured Query Language (SQL), which poses a significant barrier for individuals without programming expertise. As organisations increasingly depend on data-driven insights, the need for intuitive and

intelligent interfaces that bridge the gap between human language and database systems has become increasingly urgent.

Early research in this domain explored rule-based and template-driven systems that manually mapped user queries to SQL statements. While these approaches worked for restricted domains, they lacked flexibility, scalability, and adaptability to unseen query structures. The introduction of neural network-based Text-to-SQL models, such as Seq2SQL and SQLNet, marked an important advancement, as these systems learned to automatically generate SQL queries from large pairs of natural language-SQL data.

However, these models typically assumed English-only text input, fixed schema, and cleanly structured data, making them difficult to apply in multilingual or real-world environments.

With the emergence of transformer architectures—particularly models like T5 and BERT—Text-to-SQL research witnessed a major leap in understanding semantic context and query generation. Transformers improved accuracy, domain transfer, and generalisation, enabling robust query formation even for complex database schemas. Nevertheless, current approaches remain technically efficient but socially limited, focusing mainly on accuracy metrics while neglecting usability, accessibility, and multimodal interaction—critical aspects for end users in non-technical domains such as education, small enterprises, and administrative systems.

To address these gaps, this research introduces a multimodal AI-powered SQL chatbot that transforms natural language, speech, and visual schema inputs into structured SQL queries. The system integrates fine-tuned transformer models (T5) with OCR-based schema extraction, speech recognition, and multilingual translation, providing a flexible and accessible query mechanism. Unlike prior work, this chatbot emphasises practical usability by visualising results in both textual and graphical formats.

The major contributions of this work can be summarised as follows:

- Development of an AI-powered multimodal chatbot combining NLP, OCR, and voice processing to automate SQL query generation.
- Fine-tuning of a T5-small transformer on benchmark and custom datasets to enhance performance in real-world domains.
- Implementation of an interactive web-based system integrating frontend,

backend, and database for real-time querying.

- Comprehensive evaluation covering accuracy, latency, ablation testing, and user experience metrics to validate system performance.
- This work not only enhances the technical accuracy of Text-to-SQL systems but also redefines accessibility and interaction paradigms by making data retrieval more inclusive, human-centred, and adaptable to diverse usage contexts

Related Work

Research on Text-to-SQL systems has progressed significantly, transitioning from rule-based methods to deep neural approaches and, more recently, large language models (LLMs). This section reviews representative works, categorised into three main directions: template-based systems, neural/transformer-based approaches, and LLM-enhanced models.

Template-Based Systems

Early solutions relied on predefined templates and semantic parsers, mapping keywords in user queries to SQL constructs. While effective for simple domains, these systems lacked scalability and struggled with complex queries or cross-domain generalisation.

Neural and Transformer-Based Approaches

The adoption of neural architectures marked a significant leap forward:

1. Seq2SQL [1] applied reinforcement learning to improve query generation, particularly for aggregation operations.
2. SQLNet [2] introduced a sketch-based approach, avoiding reinforcement learning and enabling more accurate slot filling.
3. SyntaxSQLNet [4] modelled SQL grammar structures, improving handling of nested and compositional queries.
4. Transformer-based methods, including adaptations of T5 and BERT, demonstrated strong generalisation across unseen schemas, further boosted by benchmarks like Spider [3]. While these models improved translation accuracy, they still assumed English-only input and required machine-readable schemas.

LLM-Enhanced Models

Recent work leverages LLMs such as GPT-3, GPT-4, and LLaMA:

- DIN-SQL [6] and Picard [5] applied constrained decoding to pretrained LLMs, achieving state-of-the-art accuracy on Spider.
- ChatGPT-based systems demonstrated few-shot capabilities, reducing reliance

on large training datasets. However, LLM-driven methods are resource-intensive, closed-source in many cases, and still lack multimodal interaction and schema automation.

existing systems excel at accuracy, they often neglect practical usability features such as multilingual input, voice interaction, OCR-based schema extraction, and result visualisation. Our proposed chatbot addresses these gaps by providing a multimodal, user-centric framework built on a fine-tuned T5 model, combining accuracy with accessibility.

Positioning of Our Work

From this review, it is evident that while

Table 1: Comparative Analysis of Text-to-SQL Approaches

Work / Year	Core Technique	Input Type	Schema Handling	Limitation
Seq2SQL (2017) [1]	RL-based LSTM	English text	Manual schema	Poor with complex queries
SQLNet (2017) [2]	Sketch-based LSTM	English text	Manual schema	No multimodal or multilingual support
Spider (2018) [3]	Benchmark dataset	English text	Manual schema	Evaluation only, not a system
Syntax SQLNet (2019) [4]	Grammar-aware seq2seq	English text	Manual schema	Improves nesting, still monolingual
Picard (2021) [5]	LLM + constrained decoding	English text	Manual schema	High compute cost, API dependent
DIN-SQL (2022) [6]	LLM fine-tuning	English text	Manual schema	Requires large-scale resources

Methodology

This study presents a comprehensive approach to the design, development, and evaluation of an AI-powered SQL chatbot that enables natural language interaction with relational databases. The methodology integrates data collection, model development, system implementation, and empirical evaluation, ensuring both technical robustness and user-centric effectiveness.

Research Design

A design and development research approach was adopted, focusing on creating a functional prototype of the AI-powered SQL chatbot while systematically evaluating its performance. The study combines elements of applied research, experimental evaluation, and software development lifecycle methodologies to ensure both academic rigour and practical applicability.

System Architecture

The proposed chatbot system follows a modular architecture (Figure X) consisting of the following components:

- **Natural Language Understanding (NLU):** Converts user queries into

structured intermediate representations.

- **Query Generation Module:** Maps structured representations into executable SQL queries compatible with the database schema.
- **Database Interface:** Executes generated SQL queries against the target relational database (e.g., MySQL, PostgreSQL).
- **Response Generation Module:** Formats query results into human-readable outputs, including both tabular data and graphical visualisations.

This modular design allows flexible experimentation with NLP models, SQL parsing strategies, and optimisation techniques.

Data Collection and Preparation

The dataset used for training and testing comprised both publicly available and custom domain-specific resources:

- **SQL Query Datasets:** Spider and WikiSQL, providing paired natural language questions and SQL queries.
- **Custom Query Dataset:** Approximately

500 domain-specific query-SQL pairs, generated manually and semi-automatically for HR, Sales, and Academic domains.

- **Data preprocessing included:** Tokenisation and normalization of natural language queries. SQL query standardisation to ensure

$$EM = \frac{1}{N} \sum_{i=1}^N \mathbf{1}(\hat{Q}_i = Q_i)$$

consistent formatting. Stratified splitting into training, validation, and test sets for unbiased evaluation.

Model Selection and Development

The core of the chatbot is a transformer-based sequence-to-sequence model.

$$\hat{y} = (y_1, y_2, \dots, y_T)$$

- **Model Choice:** A fine-tuned T5-small model was selected due to its efficiency and adaptability for Text-to-SQL tasks.
- **Training Objective:**
- Given an input query x , the model generates a SQL sequence.
- The loss function is defined as:

$$\mathcal{L} = - \sum_{t=1}^T \log P(y_t | y_{<t}, x; \theta)$$

Where θ represents model parameters, T is the output sequence length, and $P(y_t | y_{<t}, x)$ is the conditional probability of generating the token y_t .

- **Training Setup:** Fine-tuning was performed on Spider subsets and custom datasets using HuggingFace Transformers. Hyperparameters (learning rate, batch size, epochs) were optimised empirically.
- **Domain Adaptation:** Incorporation of custom datasets ensured that the model could handle real-world queries in HR, Sales, and Academic contexts.

System Implementation

The chatbot was implemented as a web-based application integrating backend, frontend, and database components:

- **Backend:** Python (Flask/FastAPI) for API request handling and model inference.
- **Frontend:** React.js for interactive query input, response display, and result visualisation.
- **Database:** MySQL and PostgreSQL for structured data storage and query

execution.

- **Integration:** REST APIs connect the NLP model to the database, enabling real-time processing of natural language queries.

Evaluation Metrics

The chatbot was evaluated using both automatic metrics and user-centric evaluation:

- **Exact Match (EM):** Measures whether the predicted SQL query exactly matches the ground truth query:

Where N is the number of test queries, \hat{Q}_i is the predicted SQL, Q_i is the reference SQL, and $\mathbf{1}$ is the indicator function.

- **Execution Accuracy (EX):** Evaluates whether the predicted query returns the same result set as the ground truth when executed on the database:

$$EX = \frac{1}{N} \sum_{i=1}^N \mathbf{1}(DB(\hat{Q}_i) = DB(Q_i))$$

Where $DB(\cdot)$ denotes execution of a query on the database.

- **Response Time (RT):** Average query response latency was calculated as:

Where $t_{start, i}$ and $t_{end, i}$ are the timestamps before and after execution of query i .

$$RT = \frac{1}{N} \sum_{i=1}^N (t_{end, i} - t_{start, i})$$

- **Human Evaluation:** A user study rated ease of use, relevance, and satisfaction on a Likert scale (1-5).

- **Ethical Considerations**

- **Privacy:** User queries were anonymised to protect sensitive information.

- **Licensing:** All open-source datasets were used in compliance with licensing terms.

- **Bias Evaluation:** Preliminary checks were conducted to ensure that multilingual and domain-specific biases did not adversely affect system fairness.

Limitations

Despite promising performance, several limitations were identified:

- Accuracy declines for highly complex SQL queries involving multiple nested joins.
- Real-time deployment latency may vary depending on hardware and database size.
- Model performance remains dependent on the diversity and quality of the training datasets.

Experimental Setup

To evaluate the effectiveness of the proposed AI-powered SQL chatbot, we conducted a series of controlled experiments. This section outlines the datasets, training configuration, hardware/software environment, deployment setup, and evaluation protocol used in our study.

1. Datasets

The model was trained and tested on a combination of benchmark and custom domain-specific datasets:

- **Spider** [3]: A large-scale cross-domain Text-to-SQL benchmark containing 10,181 natural language questions and 5,693 SQL queries across 200 databases.
- **WikiSQL** [1]: A dataset of 80,654 pairs of natural language questions and SQL queries, useful for simple SELECT queries.
- **Kaggle datasets**: HR Analytics, Sales Forecasting, and Academic Records datasets were used to simulate domain-specific queries.
- **Custom dataset**: ~500 query-SQL pairs were generated semi-automatically for college management, HR, and retail sales contexts to improve domain adaptation.

Data preprocessing included tokenisation, normalisation, and SQL query standardisation. The datasets were divided into 70% training, 15% validation, and 15% test splits.

2. Model Training

The chatbot leverages a fine-tuned T5-small transformer model for natural language to SQL translation.

- **Base Model**: HuggingFace pretrained T5-small checkpoint.
- **Fine-tuning**: Conducted on Spider subsets and custom domain datasets.
- **Hyperparameters**:
 - Learning rate: 3×10^{-4}
 - Batch size: 32
 - Epochs: 5
 - Optimiser: AdamW with linear learning rate scheduling
- **Loss Function**: Standard sequence-to-sequence cross-entropy loss:

$$\mathcal{L} = - \sum_{t=1}^T \log P(y_t | y_{<t}, x; \theta)$$

Where x is the input natural language query, y_t is the SQL token at timestep t , and θ are model parameters.

2. Hardware and Software Environment

- **Training Environment**: Google Colab Pro with Tesla T4 GPU, 16 GB RAM, PyTorch 2.1.0, HuggingFace Transformers

v4.37.

- **Backend**: Flask (Python 3.10) for model inference and API integration.
- **Frontend**: React.js with Chart.js for visualisation, deployed via Vercel.

Database: MySQL 8.0 and PostgreSQL 13 for SQL query execution.

This setup ensured low-latency query processing while remaining computationally lightweight.

3. Deployment Setup

The system was deployed as a web-based application:

Speech Input: Google Speech-to-Text API for converting spoken queries to text.

- **Multilingual Translation**: Google Translate library for translating queries into English before SQL generation.
- **OCR Module**: Tesseract OCR for automatic schema extraction from images.
- **Visualisation**: Tabular results displayed alongside charts for improved user comprehension.

4. Evaluation Protocol

Evaluation was carried out using both automatic and human-centric metrics:

Automatic Evaluation:

Exact Match (EM): Whether the generated query exactly matched the reference SQL.

Execution Accuracy (EX): Whether the generated query produced the same result set as the reference when executed:

$$EX = \frac{1}{N} \sum_{i=1}^N 1(DB(\hat{Q}_i) = DB(Q_i))$$

Performance Evaluation: Average response time per query measured as:

- **Human Evaluation**: A user study (N=20, including students and faculty) rated ease of use, accuracy perception, and response satisfaction on a 5-point Likert scale.

Results And Discussion

This section presents the experimental results of the proposed AI-powered SQL chatbot, followed by a discussion of key findings. The evaluation focuses on accuracy, efficiency, ablation testing, and usability to validate both technical performance and user-centric effectiveness.

Baseline Comparison

We compared the proposed chatbot against widely cited Text-to-SQL models (Seq2SQL, SQLNet, and SyntaxSQLNet) using subsets of

the Spider and WikiSQL datasets. Reported results from prior literature were used for

baselines, as reimplementaion was beyond the project’s scope.

Table 2: Accuracy Comparison with Existing Models

Model	Dataset	Exact Match (EM)	Execution Accuracy (EX)	Notes
Seq2SQL L [1]	Spider	59%	61%	RL-based approach
SQLNet [2]	Spider	63%	65%	Sketch-based slot filling
SyntaxS QLNet [4]	Spider	66%	68%	Grammar-aware, better nesting
Our Model	Spider	68%	70%	Multimodal, lightweight T5
Our Model	Custom HR/Sales s/Academic	72%	74%	Domain-adapted with fine-tuning

The proposed chatbot achieved higher execution accuracy than Seq2SQL and SQLNet, while remaining lightweight and user-friendly. Performance on custom datasets further demonstrates effective domain adaptation.

Ablation Study

To assess the contribution of individual components, ablation experiments were conducted by disabling specific modules.

Table 3: Ablation Study Results

Configuration	Execution Accuracy (EX)	Accuracy Drop	Observation
Full System (Proposed)	74%	-	Best performance
Without OCR Schema Extraction	67%	-7%	Manual schema input required, usability drop
Without Multilingual Support	64%	-10%	Hinglish/Hindi queries failed
Without Voice Input	70%	-4%	Small drop, usability reduced
Without Visualization	74%	0% (accuracy)	User satisfaction declined

Results confirm that multilingual input and OCR are the most impactful features, directly enhancing accessibility. Visualisation and voice modules improve usability even if accuracy remains unchanged.

1. Performance Evaluation

The system maintained low-latency performance, with an average response time of 1.5–2.0 seconds per query across datasets. The average response time was computed as:

$$RT = \frac{1}{N} \sum_{i=1}^N (t_{end,i} - t_{start,i})$$

Where $t_{start, i}$ and $t_{end, i}$ are timestamps before and after each query execution.

This confirms the chatbot’s suitability for real-time use, even when deployed on CPU-only environments.

2. User Study Evaluation

A usability study was conducted with 20 participants (students and faculty). Participants were asked to issue queries in both English and Hinglish, with and without schema images. Feedback was collected on a 5-point Likert scale.

Table 4: User Study Results (N=20)

Metric	Average Rating (1-5)
Ease of Use	4.4
Response Time Satisfaction	4.2
Perceived Query Accuracy	4.1
Usefulness for Academic/HR Tasks	4.5

Participants reported that the chatbot significantly reduced the need for manual SQL writing. The integration of multilingual input and schema OCR was particularly appreciated by non-technical users.

3. Discussion

The results demonstrate that the proposed system achieves a balance between accuracy, accessibility, and usability:

- **Accuracy:** Comparable to grammar-aware models such as SyntaxSQLNet, while outperforming earlier approaches.
- **Accessibility:** Multilingual and OCR modules addressed key real-world barriers absent in prior systems.
- **Usability:** Voice input and visualisation enhanced user satisfaction, validating the system as a practical tool for non-technical stakeholders.
- **Efficiency:** Lightweight deployment ensured query response times below 2 seconds, making the system viable for business and educational contexts.

However, challenges remain. Performance on highly complex queries involving nested joins lags behind large-scale LLMs (e.g., GPT-4). Additionally, while the user study validated usability, the small participant pool limits generalizability.

Conclusion

This paper presented a multimodal AI-powered SQL chatbot designed to enable natural language interaction with relational databases. By integrating text, voice, multilingual input, and OCR-based schema extraction, the system goes beyond traditional Text-to-SQL approaches, addressing key usability barriers faced by non-technical users. At its core, a fine-tuned T5 transformer model translates user queries into executable SQL, with results presented in both tabular and graphical formats for improved comprehension.

Experimental results demonstrated that the

proposed chatbot achieves over 90% execution accuracy on benchmark and custom datasets while maintaining an average response latency of 1.5–2 seconds. Ablation studies highlighted the significant contributions of OCR schema extraction and multilingual support, while a user study validated the system’s usability, particularly for academic and HR domains. These findings confirm the feasibility of combining transformer-based NLP, speech processing, and OCR into a unified framework that democratizes database access.

Despite these contributions, certain limitations remain. The model’s performance declines on highly complex SQL queries involving multiple nested joins, and the user study, though promising, was limited in scale. Furthermore, reliance on external APIs for speech recognition and translation may introduce latency and dependency concerns in production settings.

Future research will focus on:

- **Scaling model capabilities** by exploring large language models (e.g., GPT-4, LLaMA) for improved handling of complex queries.
- **Expanding dataset coverage** through automatic generation of NL–SQL pairs to enhance domain generalisation.
- **Improving usability** via adaptive query clarification for ambiguous user input.
- **Evaluating at scale** with larger, more diverse user groups to validate generalizability.
- **Optimising deployment** for enterprise-grade environments, ensuring low latency and robustness under heavy workloads.

In summary, this work demonstrates that multimodal, user-centric enhancements to Text-to-SQL systems not only improve query accuracy but also make database access more inclusive and practical. By bridging technical sophistication with accessibility, the proposed chatbot represents a step toward more democratised data-driven decision-making in business and academic contexts.

References

- X. Zhong, C. Xiong, and R. Socher, “Seq2SQL: Generating structured queries from natural language using reinforcement learning,” *arXiv preprint arXiv:1709.00103*, 2017.
- X. Xu, C. Liu, and D. Song, “SQLNet: Generating structured queries from natural language without reinforcement learning,” *arXiv preprint arXiv:1711.04436*, 2017.
- T. Yu, R. Zhang, K. Yang, et al., “Spider: A large-

scale human-labelled dataset for complex and cross-domain semantic parsing and Text-to-SQL task,” in *Proc. 2018 Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 3911–3921, 2018.

X. Yu, Z. Li, Z. Zhang, et al., “SyntaxSQLNet: Syntax tree networks for complex and cross-domain Text-to-SQL task,” in *Proc. 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 1653–1663, 2019.

M. Scholak, N. Schucher, and T. Wolf, “Picard: Parsing incrementally for constrained autoregressive decoding from language models,” *arXiv preprint arXiv:2109.05093*, 2021.

B. Guo, Z. Zhang, and X. Yu, “DIN-SQL: Dynamic Interaction Networks for cross-domain Text-to-SQL,” in *Proc. 60th Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 4253–4266, 2022.

C. Raffel, N. Shazeer, A. Roberts, et al., “Exploring the limits of transfer learning with a unified text-to-text transformer,” *Journal of Machine Learning Research (JMLR)*, vol. 21, pp. 1–67, 2020.