

# Semantic Captioning: Benchmark Dataset and Graph-Aware Few-Shot In-Context Learning for SQL2Text

Ali Al-Lawati      Jason Lucas      Prasenjit Mitra

The Pennsylvania State University

University Park, PA, USA

{aha112, js15710, pmitra}@psu.edu

## Abstract

Large Language Models (LLMs) have demonstrated remarkable performance in various NLP tasks, including semantic parsing, which translates natural language into formal code representations. However, the reverse process, translating code into natural language, termed *semantic captioning*, has received less attention. This task is becoming increasingly important as LLMs are integrated into platforms for code generation, security analysis, and educational purposes. In this paper, we focus on the captioning of SQL query (SQL2Text) to address the critical need for understanding and explaining SQL queries in an era where LLM-generated code poses potential security risks. We repurpose Text2SQL datasets for SQL2Text by introducing an iterative ICL prompt using GPT-4o to generate multiple additional utterances, which enhances the robustness of the datasets for the reverse task. We conduct our experiments using in-context learning (ICL) based on different sample selection methods, emphasizing smaller, more computationally efficient LLMs. Our findings demonstrate that leveraging the inherent graph properties of SQL for ICL sample selection significantly outperforms random selection by up to 39% on BLEU score and provides better results than alternative methods. Dataset and codes are published online<sup>1</sup>.

## 1 Introduction

Large Language Models (LLMs) have advanced significantly, exhibiting paradigm-shifting capabilities in natural language processing (NLP). Notably, in-context learning (ICL) (Zhao et al., 2023) and semantic reasoning (Tang et al., 2023) abilities allow LLMs to excel in natural language understanding (NLU) and generation tasks without extensive supervised learning. This has enabled LLMs, such

<sup>1</sup>Code/Dataset Repository: <https://github.com/alwister/ast-icl>




Figure 1: Graph AST-ICL Framework for SQL2Text generation showing the pipeline from SQL input through Graph-AST methods and ICL approaches to final text output using various LLM generators (e.g., Mistral).

as ChatGPT (OpenAI, 2024), to demonstrate remarkable capabilities in generating correct code without additional training or fine-tuning. As a result, LLMs have been widely adopted in the programming community as code explainers and commenters (Chen et al., 2023a). However, there is a notable gap in the literature regarding how well these LLMs understand code.

Semantic parsing (Text2SQL), a popular downstream task has been extensively used to assess the capacity of language models to generate correct formal code segments from corresponding natural language text, e.g., using business-related questions to generate Structured Query Language (SQL) codes (Berant and Liang, 2014; Li et al., 2024). However, there is a gap in the evaluation methods and datasets for assessing generated codes. In particular, SQL2Text, which can be seen as a specific form of code summarization or translation. While Text2SQL is well-developed and studied using LLMs (Li et al., 2024; Zhang et al., 2024), SQL2Text remains underexplored, particularly in the ICL setting. This is due to the diversity of code domains and the limited availability of suitable datasets. This leads to a natural question: if LLMs are good at translating natural language to code, how well do they perform on the reverse task?

Therefore, we consider the domain of formal code and introduce a new task called *semantic captioning* as the reverse operation of semantic pars-

ing. Semantic captioning can be defined as an SQL2Text operation, where, given an SQL query, an LLM translates the formal code into a corresponding natural language question (i.e., utterance or caption). We focus on SQL code understanding because SQL is the most commonly used language for data extraction, ranking # 1 on IEEE Spectrum’s 2023 “Jobs List” (Zhang et al., 2024). Also, SQL was recently reported among the 8 most popular programming languages in 2024 (TIOBE Software BV, 2024), underscoring its widespread importance across various technology and business roles. However, the literature mainly focuses on advancing NLU in popular high-resource programming languages, such as Python (Nam et al., 2024a).

As LLM adoption grows, bidirectional SQL-LLM translation (SQL2Text and Text2SQL) becomes crucial. This capability is vital for system security, software development, and education, such as building automated support tools for risk assessment and assistive learning. Therefore, as businesses increasingly integrate AI-powered applications and automate tasks using LLMs, the ability to generate accurate and meaningful descriptions of SQL queries becomes increasingly relevant.

Thus, to investigate these gaps in knowledge, we formulate two research questions (RQs) to study semantic captioning:

**RQ1: Can we leverage existing Text2SQL dataset for semantic captioning?** We develop the first semantic captioning dataset using iterative ICL prompt-based techniques.

**RQ2: Can leveraging SQL’s graph properties for sample selection improve ICL performance?** We propose a graph-aware method and compare it to zero-shot and few-shot approaches, and we investigate Smaller LLM (e.g., Mistral) vs. Larger LLM (GPT-4) performances across all methods.

First, we use GPT-4o and vetted Text2SQL corpora to generate the first dataset for semantic captioning using a novel iterative ICL prompt on three benchmark datasets: CoSQL (Yu et al., 2019), Spider (Yu et al., 2018), and SPaC (Yu et al., 2020). Next, we perform a human expert evaluation on the generated and original datasets to maintain quality. Finally, we experiment with three smaller-sized LLMs: GPT-J-6B (Wang and Komatsuzaki, 2021), Mistral-7B (Jiang et al., 2023), and CodeLlama-7B (Roziere et al., 2023); and one larger LLM: GPT-4 (Achiam et al., 2023). In the ICL setting, we leverage the graph properties of SQL using Graph

Neural Network (GNN) and graph pooling of the SQL abstract syntax tree (AST) to select relevant prompt demonstrations. To assess our models’ performance, we utilize the vetted datasets using four evaluation metrics: BLEU-4 (Papineni et al., 2002), two variations of BERTScore (Zhang et al., 2019), and AlignScore (Zha et al., 2023) to assess the quality of LLM-generated semantic captions.

Our contributions include: (1) introducing semantic captioning, a new task focusing on SQL code understanding; (2) contributing three benchmark datasets that have been repurposed from existing Text2SQL benchmarks (Spider, CoSQL, and SparC), processed using an iterative generation for additional utterances, and evaluated using human experts; (3) proposing a novel graph-based ICL method that enhances the performance of smaller-sized LLMs; and (4) conducting experiments comparing zero-shot, few-shot (random and BM25) with our method, as well as comparing smaller vs. larger LLMs. Our findings advance the understanding of LLMs’ ability to comprehend SQL queries and generate corresponding natural language questions, facilitating the integration of LLMs into platforms, with implications for improving code security, automating documentation, and enhancing educational tools.

## 2 Related Work

**Prompt Selection** Semantic Parsing, including Text2SQL, is a popular downstream task for LLMs (Zhang et al., 2020), which has been extensively evaluated using zero-shot, multi-shot, and based on other approaches such as fine-tuning (Chang et al., 2024), and instruction-tuning (Sun et al., 2023). Recently, ICL has shown the capacity to train LLMs using prompts by providing a few demonstration examples (Brown et al., 2020). However, LLMs are sensitive to demonstration examples, which has led researchers to investigate new demonstration retrieval methods such as word overlap (Askari et al., 2023).

Several works (Min et al., 2022; Xie et al., 2022; Saunshi et al., 2021) attempt to understand why ICL works and gauge its effectiveness to motivate the selection of best fitting demonstrations (Li et al., 2023; Luo et al., 2023). Qin et al. (2023) utilize the LLM to compare the reasoning path of training demonstrations with test examples. Khalifa et al. (2023) propose an ensemble approach that combines multiple methods and averages the pre-

Language Model	Method	CoSQL-S2T				SparC-S2T				Spider-S2T			
		BERTScore-1	BERTScore-2	BLEU	AlignScore	BERTScore-1	BERTScore-2	BLEU	AlignScore	BERTScore-1	BERTScore-2	BLEU	AlignScore
GPT-J	Zero-shot	0.5831	0.4894	0.0061	0.4585	0.6158	0.5009	0.0084	0.4885	0.6177	0.497	0.0096	0.5419
	Random-2	0.7837	0.7392	0.2021	0.4943	0.8128	0.7619	0.2472	0.5346	0.8207	0.7778	0.2561	0.6036
	BM25-2	0.7779	0.7421	0.2197	0.5189	0.833 <sup>†</sup>	0.7845 <sup>†</sup>	0.289	0.5954	0.8271	0.7916	0.2799	0.6428
	AST-ICL-TOP-2	<b>0.8153<sup>†‡</sup></b>	<b>0.7783<sup>†‡</sup></b>	<b>0.2813<sup>†‡</sup></b>	<b>0.6357<sup>†‡</sup></b>	<b>0.8484<sup>†</sup></b>	<b>0.8151<sup>†‡</sup></b>	<b>0.3093<sup>†</sup></b>	<b>0.6744<sup>†‡</sup></b>	<b>0.8552<sup>†‡</sup></b>	<b>0.8261<sup>†‡</sup></b>	<b>0.3301<sup>†</sup></b>	<b>0.7201<sup>†‡</sup></b>
	AST-ICL-2	0.8082 <sup>†‡</sup>	0.7711 <sup>†‡</sup>	0.2376	0.5583	0.8037 <sup>‡</sup>	0.7635	0.2294 <sup>‡</sup>	0.5254 <sup>‡</sup>	0.8312	0.7856	0.2741	0.6327
Mistral	Zero-shot	0.7442	0.6669	0.2458	0.7464	0.7822	0.6878	0.2391	0.7787	0.7597	0.6717	0.2226	0.7717
	Random-2	0.8515	0.8145	0.3509	0.7335	0.8751	0.8351	0.376	0.764	0.875	0.8436	0.4138	0.8066
	BM25-2	0.8498	0.8143	0.3415	0.7499	0.8797	0.8423	0.3928	0.8219 <sup>†</sup>	0.8749	0.8442	0.4027	0.8275
	AST-ICL-TOP-2	<b>0.8632</b>	<b>0.8306<sup>†‡</sup></b>	<b>0.3732</b>	<b>0.8103<sup>†‡</sup></b>	<b>0.8883</b>	<b>0.8572<sup>†</sup></b>	<b>0.3977</b>	<b>0.8259<sup>†</sup></b>	<b>0.892<sup>†‡</sup></b>	<b>0.8649<sup>†‡</sup></b>	<b>0.4348</b>	<b>0.8645<sup>†</sup></b>
	AST-ICL-2	0.8577	0.82	0.3369	0.7382	0.8707	0.8386	0.3727	0.7957	0.8801	0.8462	0.3861	0.812
CodeLlama	Zero-shot	0.7285	0.674	0.1704	0.6718	0.7623	0.6979	0.1871	0.6766	0.747	0.681	0.1716	0.6985
	Random-2	0.8384	0.8019	0.2996	0.6756	0.8644	0.8269	0.3535	0.7067	0.8618	0.8318	0.353	0.7763
	BM25-2	0.835	0.8022	0.3314	0.7153	0.8715	0.832	0.3841	0.7585	0.8732	0.8471	<b>0.3769</b>	0.7709
	AST-ICL-TOP-2	<b>0.8571<sup>†‡</sup></b>	<b>0.8217<sup>†‡</sup></b>	<b>0.3578<sup>†</sup></b>	<b>0.7766<sup>†‡</sup></b>	<b>0.8818<sup>†</sup></b>	<b>0.8502<sup>†‡</sup></b>	<b>0.3945</b>	<b>0.799<sup>†</sup></b>	<b>0.878<sup>†</sup></b>	<b>0.8509<sup>†</sup></b>	0.3725	<b>0.8396<sup>†‡</sup></b>
	AST-ICL-2	0.8451	0.8085	0.3218	0.7031	0.8692	0.8348	0.3601	0.7567	0.8733	0.8413	0.3615	0.8126

Table 1: Test Results for three benchmark datasets on three LLMs using different sample selection methods for ICL. The number after the hyphen refers to the number of demonstration samples used in the prompt, which is two in this experiment. <sup>†</sup> denotes t-test statistical significance compared to Random. <sup>‡</sup> denotes t-test statistical significance compared to BM25.

dictions. Other techniques have relied on the similarity between the test sample and the demonstrations using simple text or embedding-based comparisons (Luo et al., 2023). Recently, Chen et al. (2023b) have suggested asking the LLM to provide tailored demonstrations for a given query. However, this inefficient approach may result in higher costs as it relies on a powerful LLM to suggest effective samples.

On the contrary, there is very limited research on the reverse task: semantic captioning. A few researchers have modeled the SQL2Text problem as a seq-2-seq model (Câmara et al., 2024). Xu et al. (2018) propose a seq-2-seq model that processes the SQL using GNNs. This improved the task compared to a simpler joint training of a SQL/-Text pair. However, there is a gap in the literature on a comprehensive evaluation of LLMs on this task.

**ICL Commentation** LLMs are increasingly used for code generation, and there has been a growing interest in their capacity to understand code (Nam et al., 2024b). Geng et al. (2024) study how well LLMs can provide comments based on the number of demonstrations, whereas we consider how to select the most effective demonstrations. On the other hand, Gao et al. (2023) perform an experimental study on the effect of selection order, and the number of demonstrations on the outcome for different code tasks.

**Iterative Prompt Engineering** Recent LLM advancements have introduced new ICL capabilities, shifting generative prompt-based approaches from one-step to multi-step (iterative) pipelines. This transition addresses LLMs’ generative quality issues like hallucinations and overlooked details (Adams et al., 2023; Zhang et al., 2023). Notable

innovations include the Chain of Density (CoD) method by Adams et al. (2023), producing detailed, entity-centric summaries with improved abstractiveness and reduced bias; an iterative single-agent approach for summary generation and refinement by Zhang et al. (2023); and a tri-agent LLM-based method enhancing customization by Xiao et al. (2023). However, these techniques have not been applied to improve Text2SQL and SQL2Text generation. Our work fills this gap by leveraging an iterative generation technique to produce high-quality semantic captions for SQL queries.

### 3 Problem Definition

In this section, we describe and define our iterative prompt-based generation task, which leverages LLMs’ advanced capabilities, and our graph-aware semantic captioning, which improves LLMs’ generative capabilities. Figure 3 shows an overview of our methodological framework.

#### 3.1 Iterative Generation

First, we define our iterative prompt-based generative task that uses ICL and iterative generation to produce high-quality semantic captions. To gauge the performance of the SQL2Text problem, it is imperative to have datasets that support this problem. Existing Text2SQL datasets fail to cater well to this problem since a SQL query may be described in several different ways using natural language. In order to repurpose these datasets, we investigate an iterative process-based approach to generate and refine multiple utterances for a given query. Next, we verify that this approach consistently affects the test results of ICL. It is important to note that simple paraphrasing and one-step ICL generation such as SumCoT (Wang et al., 2023) are prone to

inconsistencies.

**Definition 1.** We define our iterative prompt-based ICL approach for generating high-quality semantic captions as a three-step process. Let  $x$  be an input SQL query. Our approach consists of:

- **Generation:** Produce a set  $\mathbf{U} = \{u_1, u_2, u_3\}$  of three distinct utterances or questions accurately representing the input SQL query  $x$ .
- **Feedback:** Generate a feedback set  $\mathbf{F} = \{f_1, f_2, f_3\}$  corresponding to each utterance in  $\mathbf{U}$ . Each  $f_i$  identifies errors, ambiguities, or areas for improvement based on four criteria: accuracy, clarity, fluency, and relevance to the original SQL query  $x$ .
- **Refinement:** Produce a refined set of utterances  $\mathbf{U}' = \{u'_1, u'_2, u'_3\}$  by applying the feedback in  $\mathbf{F}$  to the original set  $\mathbf{U}$ , optimizing for accuracy, clarity, fluency, and relevance.

This iterative process aims to maximize the quality of the semantic captions while ensuring an LLM can accurately use them for the SQL2Text generation task. Given an inference LLM,  $G$ , the refined utterances  $u'_i \in \mathbf{U}'$  should reconstruct the original SQL query  $x$  when used as input to  $G$ .  $G(u'_i)$  should yield  $x'$ , which is semantically equivalent to  $x$ .

### 3.2 Semantic Captioning

Second, we define our proposed method, leveraging graph properties for few-shot sampling to improve LLMs’ semantic captioning capabilities using SQL2Text. We experiment with multiple selection techniques to provide ICL with a better set of demonstration samples. In particular, SQL can be naturally represented as an AST: a graph-based knowledge structure. First, we assess whether a graph embedding method can best guide the sample selection problem and enhance performance compared to random or word similarity selection approaches such as BM25 (Robertson et al., 2009). Then, we benchmark the performance improvements against zero-shot. Next, we measure how well our GNN method helps small models compared to large models such as GPT-4. Finally, we verify whether any improvement is garnered by applying our method to a large model.

**Definition 2.** We define semantic captioning with prompt-based few-shot ICL for text generation as follows. Given a demonstration dataset  $\mathcal{M}$  containing pairs of SQL queries and corresponding

natural language question examples  $(x, y)$ , we aim to construct an effective prefix input prompt ( $\mathcal{P}$ ). Our prefix prompt template comprises three parameters  $\mathcal{P}(t + e + s)$ : instructional text ( $t$ ), ICL SQL examples ( $e$ ), and an SQL seed ( $s$ ).

We utilize a knowledge-aware representation retriever  $R$  to select appropriate SQL query ICL example pairs  $\mathbf{N} = \{(x_1, y_1), \dots, (x_n, y_n)\}$  from the dataset  $\mathcal{M}$  to generate a text sequence  $y'$  in the form of a question or utterance such that the similarity metric  $\mathcal{D}(y, y')$  is maximized.

Given an inference LLM,  $G$ , a good input prompt should lead to the target output sequence where the test example  $x_{\text{test}}$  is the SQL seed,  $s$ , concatenated to the prompt  $\mathcal{P}$  and passed as a prefix prompt to  $G$ . Specifically, the decoding from the LLM  $G(\mathcal{P}; x_{\text{test}})$  should yield  $y_{\text{test}}$ .

### 3.3 LLMs

We employ LLMs for dataset generation and semantic captioning evaluation. For dataset generation, we employ OpenAI’s GPT-4o<sup>2</sup> to create utterances from SQL queries in Text2SQL corpora. This LLM was selected for its proficiency in programming and natural language tasks.

For evaluation, we use: (1) smaller open-source LLMs: GPT-J-6B, Mistral-7B, and CodeLlama-7B; and (2) GPT-4 as a larger LLM for comparative evaluation. All models except GPT-4 are accessible via Hugging Face. Models details are provided in Appendix A.3.

### 3.4 Datasets

There is no existing dataset for this task, therefore we explore benchmark Text2SQL datasets. First, we verify that these Text2SQL datasets are applicable to our task. We vetted over 8 benchmark Text2SQL datasets. We verify their appropriateness by ensuring that each dataset has SQL and utterance question pairs and that each question/utterance pair does not contain schema elements such as tables or columns to avoid inconsistency, which can lead to leakage, providing additional context of the query schema that may make it easier for the language model to resolve. The author manually checked each sample across these datasets for inconsistencies. Finally, we selected the columns containing the text pairs (SQL and associated utterance), to repurpose them for the SQL2Text task. We selected three benchmarks Text2SQL datasets:

<sup>2</sup><https://openai.com/index/hello-gpt-4o/>

CoSQL, Spider, and SParC as shown in Table 5, for our proposed task.

## 4 RQ1: Iterative ICL Generation

To create the semantic captioning dataset, we processed the test sets of the selected datasets. For CoSQL, we leveraged the entire test set of 293, while for the other datasets, we randomly sampled 293 instances, totaling 879 samples for our generation task. We utilized OpenAI’s recently released GPT-4o model for utterance generation. Through prompt engineering, we created a prompt template with a JSON structure that guides the LLM through three iterative processes: generate, review, and refine. Our iterative approach is similar to the CoD iterative generation method but employs feedback based on predefined criteria—*accuracy*, *clarity*, *fluency*, and *relevance*—following the dimensions defined by Peyrard (2019). This feedback helps refine the initial generation using fewer steps. Additionally, our predefined JSON structure enhances generative controllability. The prompt consists of instructions to guide the LLM through the generation task (see Appendix A.10).

### 4.1 Evaluation Metrics

To validate the quality of our generated utterances, we assess both semantic and logical consistency to detect any hallucinations and semantic discrepancies. Following the F3 method proposed by Lucas et al. (2023), we utilize (1) AlignScore for logical alignment and (2) BERTScore for semantic alignment.

AlignScore (0-1) assesses logical consistency between the original utterance of the SQL query and the generated one, with higher scores indicating better alignment. BERTScore (0-1) evaluates semantic similarity to a reference, with higher scores reflecting greater accuracy. Low scores in either metric suggest potential inconsistencies or inaccuracies. Details of these metrics are provided in Appendix A.2.

### 4.2 RQ1 Results

We generated a dataset of 2637 utterances; each of the 879 original SQL queries has a group of three utterances. We conducted an evaluation process comprising both automated evaluation and human assessment to ensure robust quality validation of semantic captioning.

Dataset	SQL Query	Original Utterance	Generated Utterance
CoSQL-S2T	290	281	870
SParC-S2T	289	282	867
Spider-S2T	274	270	822

Table 2: Dataset Statistics post quality check. The new datasets were renamed with a ‘-S2T’ postfix for clarity which abbreviates SQL2Text.

#### 4.2.1 Generation Quality Check

**Automatic Evaluation** For each dataset, we employed an iterative prompt-based approach that generates six references per SQL query: three initial and three final references. We compare the original human written utterance with the generated initial utterance (iterative prompt, step 1) and refined final utterance (iterative prompt, step 3) using AlignScore and BERTScore. Our iterative prompt-based approach improved our average final AlignScore by 7%. Similarly, we measure semantic consistency, showing BERTScore greater than 0.85 on average. We selected the top three references based on Alignscore and BERTScore to ensure our generated utterances maintain logical and semantic consistency with gold-label references.

**Human Expert Evaluation** To ensure a more reliable technical assessment of the generated data, we prioritized domain expertise over crowdsourcing by engaging three evaluators: two authors and one external expert with extensive NLP and database experience. The evaluators independently assessed a total of 879 queries across four datasets using a structured evaluation framework with the following criteria: (1) *Semantic Accuracy*: Evaluating the correctness of natural language interpretation (2) *SQL Logic Preservation*: Assessing retention of query constraints and operations and (3) *Technical Precision*: Assessing accurate representation of SQL operations.

For each sample, we conducted a quality assessment using the SQL query, gold-label reference, and our selected top three generated references. The evaluation followed a four-point scale: (1) *Full Inconsistency*: Generated samples are considered inconsistent when any references incorrectly reflect (semantically, logically, and technically) all key elements of both the gold label and query, (2) *SQL Inconsistency*: Generated samples are deemed inconsistent if one or more references fail to correctly reflect (semantically, logically, and technically) all key elements of the query (3) *Gold-label Inconsistency*: Generated samples are marked inconsistent

if one or more references do not correctly reflect (semantically and logically) all key elements of the gold-label, and (4) *Full Consistency*: Generated samples are considered consistent when all references correctly reflect (semantically, logically, and technically) all key elements of both the gold label and query.

**Results and Validation** We removed all inconsistent original reference utterances and SQL queries and generated captions to the SQL query, resulting in a final dataset of 853 SQL queries, 832 original utterances, and 2559 generated utterances (See Table 2). The evaluation revealed strong inter-rater reliability (Fleiss’ Kappa  $\kappa = 0.802$ ) and demonstrated that our iterative ICL prompt generates high-quality captions. This evaluation framework validates the quality of our generated dataset for advancing research in the field.

## 5 RQ2: Graph-aware ICL Few-shot

We now describe our proposed method and investigate a graph-aware ICL sample selection approach on the SQL2Text task. We developed two approaches for this task using the AST. Our first method, AST-ICL, leverages clustering with random selection from the cluster, and our second, AST-ICL-TOP, is a variant that selects the top  $n$  similar examples. We compare our method to the following baselines (1) a random selection method, i.e., demonstration samples are selected using chance from the training set, and (2) a word overlap method, BM25.

### 5.1 AST-ICL

**AST** Our intuition for building a graph from the SQL for sample selection is based on the premise that the structure of the query (complex clauses, table joins, and nesting) is more informative than the word overlap. Figure 2 depicts the breakdown of an SQL query into a graph structure based on its AST. The AST provides a hierarchical structure that represents the syntactic structure of the SQL statement, such as keywords, operators, and operands.

**Graph Construction** Next, each node in the constructed AST graph is tokenized and embedded. This is followed by a Graph Convolution Network (GCN) aggregation steps:

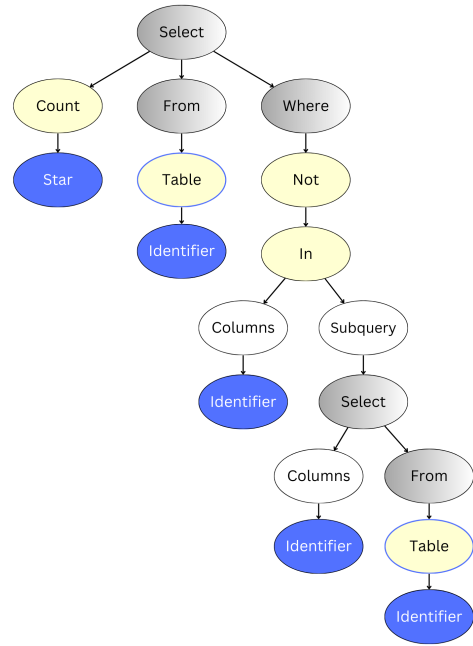


Figure 2: Parse Tree for SQL Query: SELECT count(\*) FROM Dogs WHERE dog\_id NOT IN (SELECT dog\_id FROM Treatments)

$$\mathbf{h}_i^{(l+1)} = \sigma \left( \sum_{j \in \mathcal{N}(i)} \frac{1}{\sqrt{d_i d_j}} \mathbf{W}^{(l)} \mathbf{h}_j^{(l)} \right) \quad (1)$$

where  $\mathcal{N}(i)$  is the set of neighbors of node  $i$ ,  $d_i$  and  $d_j$  are the degrees of nodes  $i$  and  $j$ ,  $\mathbf{W}^{(l)}$  is the weight matrix of layer  $l$ , and  $\sigma$  is a non-linear activation function (e.g., ReLU).

To obtain a single vector representation for each AST, we perform graph pooling:

$$\mathbf{h}_{\text{graph}} = \frac{1}{|V|} \sum_{i \in V} \mathbf{h}_i^{(L)} \quad (2)$$

where  $|V|$  is the number of nodes in the graph and  $\mathbf{h}_i^{(L)}$  is the node embedding at the final GCN layer  $L$ .

The result of the graph pooling operation is a fixed-size embedding vector  $\mathbf{h}_{\text{graph}}$  for the entire AST, representing the SQL query in latent space.

Finally, we apply clustering to the graph embeddings to group similar SQL queries. During testing, the demonstration examples are picked randomly from the corresponding cluster. This improves the time complexity of the sample selection by statistical sampling from the corresponding cluster of the test query. We refer to the variation as AST-ICL (see Figure 3).




Figure 3: AST-ICL Model Overview: First, AST graphs are constructed for each SQL segment, then encoded into vectors by a GNN embedding model, and finally clustered using k-means for AST-based In-Context Learning (ICL).

**AST-ICL-TOP** This variation follows the same approach of AST-ICL to generate the SQL AST-based graphs and embeddings. However, a full scan of the training set is used to select the top  $n$  similar samples. The computational requirement of this method is slightly higher since a full scan of the training set is performed to select the required samples instead of random selection from the corresponding cluster.

## 5.2 Baselines

We utilize two baselines: random selection, and BM25 word similarity. These two baselines establish a comprehensive evaluation framework to assess the effectiveness of our proposed AST-ICL sample selection method for ICL in SQL2Text tasks. The baselines are described in more detail in Appendix A.4.

## 5.3 Evaluation Metrics

To assess the performance of the AST-ICL in comparison to the baselines, we utilize three benchmark metrics: BERTScore, BLEU-4, and AlignScore. BERTScore measures the similarity between the generated text in latent space. For BERTScore, we employ two fine-tuned transformer models optimized for paraphrasing: paraphrase-MiniLM-L6-v2 (BERTScore-1) and paraphrase-distilroberta-base-v1 (BERTScore-2) (Reimers and Gurevych, 2019).

BLEU-4 calculates the geometric mean of n-gram precision scores (up to 4-grams) between generated text and reference, with a brevity penalty. It evaluates SQL-to-English translation quality by measuring n-gram overlap between generated text and reference questions.

AlignScore measures the logical consistency between the generated text and the reference utter-

ances by evaluating the alignment of semantic content and reasoning patterns. Details of these metrics are provided in Appendix A.2.

## 5.4 RQ2 Results

First, we seek to understand whether our proposed method can outperform baseline few-shot technique. Second, we investigate LLM zero-shot capabilities. Third, we aim to understand how well open-source small-parameter LLMs compare to larger proprietary LLMs.

### 5.4.1 AST-ICL vs. ICL few-shot

We investigate the performance of our proposed method compared to the baseline few-shot methods. We examine the ability of AST-ICL to improve LLM semantic captioning capabilities and compare it to the current SOTA method.

**Our Graph-based method outperforms baseline ICL’s few-shot sample selection for SQL semantic captioning.** Table 1 presents the results for the benchmark datasets on the metrics. This experiment only used two demonstration samples for each method. This explains the higher-than-expected performance of random selection on some of the results. However, it is easily observable that AST-ICL-TOP achieves significantly higher performance than all other methods across all metrics.

**Additional samples significantly improve the SQL2Text task performance of AST-ICL methods.** In the AST-ICL-TOP experiment, we measure the effect of adding additional samples on each method on the results. This is tabulated in Table 9 for the CoSQL-S2T dataset for brevity. Additional results are reported in Appendix A.9. Other methods also observe improvement, but to a lesser extent with random selection. These findings suggest that some samples may have limited positive influence

or are redundant.

Lang. Model	Method	CoSQL-S2T Dataset		
		BERTScore-1	BERTScore-2	BLEU
GPT-4	Zero-Shot	0.8944	0.8724	0.4649
GPT-J	BM25-4	0.8101	0.7767	0.259
	<b>AST-ICL-TOP-4</b>	0.8333	0.7974	0.2964
	AST-ICL-4	0.8171	0.7778	0.2601
	BM25-8	0.824	0.7944	0.3
	<b>AST-ICL-TOP-8</b>	0.835	0.7994	0.3125
Mistral	AST-ICL-8	0.8294	0.7903	0.2785
	BM25-4	0.8643	0.8298	0.3674
	<b>AST-ICL-TOP-4</b>	0.8722	0.8331	0.3799
	AST-ICL-4	0.862	0.834	0.3851
	BM25-8	0.8785	0.8452	0.396
CodeLlama	<b>AST-ICL-TOP-8</b>	0.8722	0.838	0.3988
	AST-ICL-8	0.872	0.839	0.4112
	BM25-4	0.8478	0.8177	0.3566
	<b>AST-ICL-TOP-4</b>	0.8667	0.832	0.356
	AST-ICL-4	0.8557	0.8311	0.339
CodeLlama	BM25-8	0.8664	0.8382	0.3935
	<b>AST-ICL-TOP-8</b>	0.8699	0.8327	0.3709
	AST-ICL-8	0.8658	0.8343	0.3611

Table 3: Test Results for competitive small-parameter LLMs vs GPT-4 Zero-shot

#### 5.4.2 Zero-shot Vanilla v.s. Few-shot

We investigate the performance of zero-shot methods compared to baseline few-shot method. We compare the capacity of LLMs to understand and translate SQL with no demonstration sample.

**The performance of the models on zero-shot was quite poor.** A cursory review of the output suggests that some results included or consisted entirely of more SQL. Different prompt restrictions were used to avoid additional code generation. A key finding is that while LLMs tested performed poorly in zero-shot, experiments suggest that only a small number of ICL samples can improve the result (Table 1).

#### 5.4.3 Smaller vs. Larger LLM

We investigate the performance of three pre-trained smaller-sized LLMs: GPT-J, Mistral, and CodeLlama. We utilize GPT-4 as our larger-sized LLM, given its superior performance across NLP tasks. GPT-4 is reported to have 100 trillion parameters. This is about 14k times larger than the small models considered in this work. Here we attempt to measure the capacity of GPT-4 to understand SQL code, and compare the performance gains achieved from AST-ICL with GPT-4. Appendix A.5 includes the prompt used with GPT-4, which was verified using OpenAI playground to generate the intended results.

GPT-4 was shown to perform relatively well on the SQL2Text task using zero-shot. Here we inves-

tigate if the performance of GPT-4 can be improved with ICL, and whether AST-ICL would outperform a random selection in this setting (Table 4).

**GPT-4 shows some performance improvement using ICL.** This finding demonstrates that ICL do not contribute to helping large LLMs as well as smaller LLMs in improving code understanding (Table 4). However, there is a notable improvement in BLEU score.

**There is no clear advantage in using AST-ICL sample selection vs. other approaches using GPT-4.** Larger models understand SQL very well in zero-shot, so adding ICL does not add any significant improvement (Table 4).

Method	CoSQL-S2T Dataset		
	BERTScore-1	BERTScore-2	BLEU
GPT-4 Zero-shot	0.8944	0.8724	0.4649
GPT-4 + BM25-2	0.9035	0.8655	0.5405
GPT-4 + <b>AST-ICL-TOP-2</b>	0.8999	0.8652	0.5107
GPT-4 + Random-2	0.9017	0.8701	0.5249

Table 4: Test Results for GPT-4 using different ICL prompts

## 6 Discussion

**AST-ICL demonstrates a remarkable capacity to improve the effectiveness of smaller LLMs for semantic captioning.** Our graph-aware methods for SQL2Text outperform alternative sample selection methods for ICL, such as BM25, while providing more efficiency. In particular, the capacity to choose more samples achieved performance slightly lower than GPT-4 from smaller models by orders of magnitude. The GNN-based encoding method of AST-ICL and AST-ICL-TOP using the syntactical elements of SQL based on the AST indicates that LLMs exhibit enhanced reasoning ability when the examples are structurally similar, compared to word similarity in most cases.

**Smaller LLMs need ICL for semantic captioning.** Even though SQL is a popular language incorporated in the training parameter of smaller LLMs, ICL is integral to achieving decent performance on SQL2Text. This suggests that smaller LLMs cannot follow prompt instructions very well for this task, unlike GPT-4, which can perform the task correctly in zero-shot.

**Random ICL quickly suffers from diminishing returns.** An ICL sample selection method based on a strategy, even a simple one, appears to perform better with more samples. However, ran-



dom selection can perform well with a few samples (e.g. 2 samples), and outperform BM25 on some metrics. This suggests that word similarity may be counter-productive for small sample sizes, and may provide limited diversity.

**AST-ICL improves GPT-4, but insignificantly.** While AST-ICL-TOP improves GPT4, our comparative analysis reveals important insights about the relationship between model size and semantic captioning performance. Table 3 shows GPT-4, with approximately 100 trillion parameters, achieves a zero-shot BLEU score of 0.4649 on the CoSQL-S2T dataset. However, smaller models enhanced with AST-ICL demonstrate competitive performance. Specifically, Mistral-7B with AST-ICL achieves a BLEU score of 0.4112, maintaining a comparative performance to GPT-4 while requiring significantly fewer computational resources. The efficiency advantages of smaller models become particularly evident in production environments. This efficiency-performance trade-off provides a compelling argument for using smaller models in practical applications.

## 7 Conclusion

Our work introduces the problem of semantic captioning as the inverse problem of the well-known semantic parsing downstream task, focusing on smaller parameter LLMs, which may be easier to integrate into coding and educational platforms. In order to motivate additional work on this problem, we have provided a method to repurpose, generate, and publish new datasets, and provide an iterative prompt for the community to reuse for additional datasets.

Furthermore, we demonstrate that for the SQL2Text task of semantic parsing, leveraging the graph properties of SQL can provide added value. These findings have the potential to facilitate the integration of smaller LLMs into coding and educational platforms, making them more accessible and efficient for a wide range of users. Moreover, our work contributes to the broader understanding of LLMs' capabilities in code comprehension and generation tasks, paving the way for further research and applications in this domain.

## 8 Limitations and Future Work

### 8.1 Current Limitations

Our work has several key limitations that warrant consideration. First, while our approach demon-

strates effectiveness with SQL, its applicability may vary across different programming languages and domains. Second, although we utilize established benchmarks, our current datasets may not fully capture the complexity and diversity of real-world SQL queries and their natural language descriptions. In our preliminary experiments with datasets like KaggleDBQA, we encountered excessive inconsistencies that necessitated excluding them from our analysis, highlighting the challenges in dataset quality and curation.

### 8.2 Future Research Directions

Several important research directions emerge from our work. First, our experience with inconsistent datasets underscores the critical need for robust quality metrics in SQL-text paired datasets. Future work should focus on systematic evaluation of existing benchmarks and creation of new, high-quality datasets that better reflect real-world complexity. The development of standardized quality assessment protocols will be essential for advancing the field.

In terms of technical advancements, as the field evolves, research should investigate the scaling behavior of AST-ICL with larger models and explore potential hybrid approaches combining different model sizes. The development of more sophisticated graph representation techniques will be crucial, as will extending our approach to other programming languages and domains. These technical improvements could significantly enhance the versatility and effectiveness of semantic captioning systems.

Additionally, the evaluation framework for semantic captioning systems requires substantial development. Future research should expand evaluation criteria to incorporate domain-specific requirements while developing standardized quality assessment protocols. Incorporating more comprehensive human evaluation methodologies will be crucial for validating system performance across diverse use cases. These enhancements to evaluation frameworks will provide more reliable insights into system performance and further advance semantic captioning.

These directions could potentially improve performance while maintaining computational efficiency. The emergence of more powerful LLMs presents opportunities to push performance boundaries further, particularly in addressing the dataset quality challenges we encountered. Additionally,

expanding the evaluation framework to include more diverse SQL queries and domain-specific applications would provide valuable insights for practical deployments.

## Acknowledgments

The authors sincerely thank Rupak Kumar Das, a PhD student, for participating in the human evaluation, which significantly contributed to the quality of this study. We also extend our gratitude to Dr. Adaku Uchendu, AI Researcher at MIT Lincoln Lab, for her important guidance and support throughout this work.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Griffin Adams, Alex Fabbri, Faisal Ladhak, Eric Lehman, and Noémie Elhadad. 2023. From sparse to dense: Gpt-4 summarization with chain of density prompting. In *Proceedings of the 4th New Frontiers in Summarization Workshop*, pages 68–74.
- Arian Askari, Amin Abolghasemi, Gabriella Pasi, Wessel Kraaij, and Suzan Verberne. 2023. Injecting the bm25 score as text improves bert-based re-rankers. In *European Conference on Information Retrieval*, pages 66–83. Springer.
- Jonathan Berant and Percy Liang. 2014. Semantic parsing via paraphrasing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1415–1425.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Vanessa Câmara, Rayol Mendonca-Neto, André Silva, and Luiz Cordovil. 2024. A large language model approach to sql-to-text generation. In *2024 IEEE International Conference on Consumer Electronics (ICCE)*, pages 1–4. IEEE.
- Chih-Yung Chang, Yuan-Lin Liang, Shih-Jung Wu, and Diptendu Sinha Roy. 2024. Sv2-sql: a text-to-sql transformation mechanism based on bert models for slot filling, value extraction, and verification. *Multi-media Systems*, 30(1):1–17.
- Eason Chen, Ray Huang, Han-Shin Chen, Yuen-Hsien Tseng, and Liang-Yi Li. 2023a. Gptutor: a chatgpt-powered programming tool for code explanation. In *International Conference on Artificial Intelligence in Education*, pages 321–327. Springer.
- Wei-Lin Chen, Cheng-Kuang Wu, Yun-Nung Chen, and Hsin-Hsi Chen. 2023b. Self-ICL: Zero-Shot In-Context Learning with Self-Generated Demonstrations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15651–15662, Singapore. Association for Computational Linguistics.
- Shuzheng Gao, Xin-Cheng Wen, Cuiyun Gao, Wenxuan Wang, Hongyu Zhang, and Michael R. Lyu. 2023. What Makes Good In-Context Demonstrations for Code Intelligence Tasks with LLMs? In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 761–773, Luxembourg, Luxembourg. IEEE.
- Mingyang Geng, Shangwen Wang, Dezun Dong, Haotian Wang, Ge Li, Zhi Jin, Xiaoguang Mao, and Xiangke Liao. 2024. Large Language Models are Few-Shot Summarizers: Multi-Intent Comment Generation via In-Context Learning. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering, ICSE '24*, pages 1–13, New York, NY, USA. Association for Computing Machinery.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Muhammad Khalifa, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, and Lu Wang. 2023. Exploring Demonstration Ensembling for In-context Learning. *arXiv preprint. ArXiv:2308.08780 [cs]*.
- Jinyang Li, Binyuan Hui, Ge Qu, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, et al. 2024. Can llm already serve as a database interface? a big bench for large-scale database grounded text-to-sqls. *Advances in Neural Information Processing Systems*, 36.
- Xiaonan Li, Kai Lv, Hang Yan, Tianyang Lin, Wei Zhu, Yuan Ni, Guotong Xie, Xiaoling Wang, and Xipeng Qiu. 2023. Unified Demonstration Retriever for In-Context Learning. *arXiv preprint. ArXiv:2305.04320 [cs]*.
- Jason Lucas, Adaku Uchendu, Michiharu Yamashita, Jooyoung Lee, Shaurya Rohatgi, and Dongwon Lee. 2023. Fighting fire with fire: The dual role of llms in crafting and detecting elusive disinformation. In *2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023*, pages 14279–14305. Association for Computational Linguistics (ACL).
- Man Luo, Xin Xu, Zhuyun Dai, Panupong Pasupat, Mehran Kazemi, Chitta Baral, Vaiva Imbrasaitė, and Vincent Y. Zhao. 2023. Dr.ICL: Demonstration Retrieved In-context Learning. *arXiv preprint. ArXiv:2305.14128 [cs]*.

- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?](#) *arXiv preprint*. ArXiv:2202.12837 [cs].
- Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. 2024a. Using an llm to help with code understanding. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, pages 1–13.
- Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. 2024b. [Using an LLM to Help With Code Understanding](#). In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering, ICSE '24*, pages 1–13, New York, NY, USA. Association for Computing Machinery.
- OpenAI. 2024. Chatgpt: Large language model. <https://openai.com/chatgpt>. Accessed: 2024-12-16.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Maxime Peyrard. 2019. A simple theoretical model of importance for summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1059–1073.
- Chengwei Qin, Aston Zhang, Anirudh Dagar, and Wenming Ye. 2023. [In-Context Learning with Iterative Demonstration Selection](#). *arXiv preprint*. ArXiv:2310.09881 [cs].
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- Nikunj Saunshi, Sadhika Malladi, and Sanjeev Arora. 2021. [A Mathematical Exploration of Why Language Models Help Solve Downstream Tasks](#). *arXiv preprint*. ArXiv:2010.03648 [cs, stat].
- Gang Sun, Ran Shen, Liangfeng Jin, Yifan Wang, Shiyu Xu, Jinpeng Chen, and Weihao Jiang. 2023. Instruction tuning text-to-sql with large language models in the power grid domain. In *Proceedings of the 2023 4th International Conference on Control, Robotics and Intelligent System*, pages 59–63.
- Xiaojuan Tang, Zilong Zheng, Jiaqi Li, Fanxu Meng, Song-Chun Zhu, Yitao Liang, and Muhan Zhang. 2023. Large language models are in-context semantic reasoners rather than symbolic reasoners. *arXiv preprint arXiv:2305.14825*.
- TIOBE Software BV. 2024. Tiobe programming community index. <https://www.tiobe.com/tiobe-index>. Accessed: 2024-12-26.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Yiming Wang, Zhuosheng Zhang, and Rui Wang. 2023. Element-aware summarization with large language models: Expert-aligned evaluation and chain-of-thought method. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Wen Xiao, Yujia Xie, Giuseppe Carenini, and Pengcheng He. 2023. Chatgpt-steered editing instructor for customization of abstractive summarization. *arXiv preprint arXiv:2305.02483*, page 103.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2022. [An Explanation of In-context Learning as Implicit Bayesian Inference](#). *arXiv preprint*. ArXiv:2111.02080 [cs].
- Kun Xu, Lingfei Wu, Zhiguo Wang, Yansong Feng, and Vadim Sheinin. 2018. Sql-to-text generation with graph-to-sequence model. *arXiv preprint arXiv:1809.05255*.
- Tao Yu, Rui Zhang, He Yang Er, Suyi Li, Eric Xue, Bo Pang, Xi Victoria Lin, Yi Chern Tan, Tianze Shi, Zihan Li, et al. 2019. Cosql: A conversational text-to-sql challenge towards cross-domain natural language interfaces to databases. In *2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019*, pages 1962–1979. Association for Computational Linguistics.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921.
- Tao Yu, Rui Zhang, Michihiro Yasunaga, Yi Chern Tan, Xi Victoria Lin, Suyi Li, Heyang Er, Irene Li, Bo Pang, Tao Chen, et al. 2020. Sparc: Cross-domain semantic parsing in context. In *57th Annual Meeting of the Association for Computational Linguistics, ACL 2019*, pages 4511–4523. Association for Computational Linguistics (ACL).
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. Alignscore: Evaluating factual consistency

with a unified alignment function. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348.

Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023. Summit: Iterative text summarization via chatgpt. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10644–10657.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

Tingkai Zhang, Chaoyu Chen, Cong Liao, Jun Wang, Xudong Zhao, Hang Yu, Jianchao Wang, Jianguo Li, and Wenhui Shi. 2024. Sqlfuse: Enhancing text-to-sql performance through comprehensive llm synergy. *arXiv preprint arXiv:2407.14568*.

Xiaoyu Zhang, Fengjing Yin, Guojie Ma, Bin Ge, and Weidong Xiao. 2020. M-sql: Multi-task representation learning for single-table text2sql generation. *IEEE Access*, 8:43156–43167.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

## A Supplement on Experimental Setting

### A.1 Description of Datasets

The SQL2Text datasets we used are as follows:

**CoSQL** (Yu et al., 2019): A large-scale dataset for cross-domain, general-purpose database querying dialogue systems, consisting of 30k+ turns with 10k+ annotated SQL queries from 3k dialogues querying 200 databases across 138 domains.

**Spider** (Yu et al., 2018): A large-scale, complex, and cross-domain semantic parsing and Text2SQL dataset with 10K+ questions and 5K+ unique complex SQL queries on 200 databases covering 138 domains, focusing on complex SQL queries (JOIN, GROUP BY, nested queries) and requiring generalization to new queries and schemas.

**SParC** (Yu et al., 2020): A dataset for cross-domain Semantic Parsing in Context, consisting of 4K+ coherent question sequences (12k+ questions with SQL queries) from controlled user interactions with 200 databases over 138 domains, demonstrating complex contextual dependencies and requiring generalization to new domains.

### A.2 Evaluation Metrics

We describe the evaluation metrics we employed for assessing the quality of generated utterances and for our semantic caption generation task

Dataset	Count	Type	Query Length			Keywords		
			min	max	mean	min	max	mean
CoSQL	2159	Train	6	87	19.7	2	14	5.5
	293	Val	6	58	19.8	2	12	5.5
SParC	3034	Train	4	87	17.6	2	14	5.0
	422	Val	4	58	17.5	2	13	5.0
Spider	7000	Train	4	87	15.9	2	14	4.6
	1034	Val	4	58	15.6	2	13	4.6

Table 5: Statistics of query length and SQL keyword occurrence (e.g. SELECT, FROM) in training and validation samples across different datasets

- **BERTScore** (Reimers and Gurevych, 2019): we employ two transformer variants fine-tuned for paraphrasing on different general datasets: (1) paraphrase-MiniLM-L6-v2<sup>3</sup> (BERTScore-1) and (2) paraphrase-distilroberta-base-v1<sup>4</sup> (BERTScore-2). We selected these fine-tuned transformer variants given their optimization for paraphrasing and their strong performance on semantic search tasks. BERTScore measures the similarity between the generated text and the original prompt or question using their corresponding transformer embeddings.
- **BLEU-4** (Papineni et al., 2002): we employ BLEU-4 (BLEU), which calculates the geometric mean of n-gram precision scores (up to 4-grams) between the generated text and the reference, multiplied by a brevity penalty to penalize short translations. BLEU evaluates the quality of the translation from SQL to English by measuring the overlap of n-grams between the generated text and the reference question.
- **AlignScore** (Zha et al., 2023): we employ AlignScore, a factual alignment or logical consistency evaluation metric designed to measure how well the generated text aligns with the original source material in terms of meaning and factual correctness. It compares factual statements in the generated text with the source to assess consistency. This metric is particularly useful for evaluating tasks like text summarization and semantic captioning while preserving original meaning and facts.

<sup>3</sup><https://huggingface.co/sentence-transformers/paraphrase-MiniLM-L6-v2>

<sup>4</sup><https://huggingface.co/sentence-transformers/paraphrase-distilroberta-base-v1>

### A.3 LLMs

We leverage several LLMs of varying sizes for two primary purposes: (1) the generation of utterances to develop a semantic caption dataset using SQL queries from benchmark Text2SQL corpora and (2) the generation of semantic captions using our proposed AST-ICL and baseline methods.

We employ the following LLM for semantic caption generation:

- **GPT-4o**: A multimodal model that accepts text or image inputs and generates text outputs. It offers comparable intelligence to GPT-4 Turbo, but with twice the generation speed and 50% lower cost. GPT-4o also provides superior vision capabilities and excels in non-English languages, making it highly efficient for tasks like semantic captioning.

We utilize the following LLMs for the evaluation of generated semantic captions:

- **GPT-J-6B** (Wang and Komatsuzaki, 2021): A GPT-based LLM developed by EleutherAI, enhanced for complex NLP generation and understanding applications<sup>5</sup>.
- **Mistral-7B** (Jiang et al., 2023): A transformer-based LLM optimized for large-scale language understanding and generation tasks<sup>6</sup>.
- **CodeLlama-7B** (Roziere et al., 2023): A custom transformer architecture developed by Meta AI, efficient and scalable for a variety of NLP tasks, with a recent focus on code understanding and generation<sup>7</sup>.
- **GPT-4** (Achiam et al., 2023): OpenAI’s advanced multimodal LLM, excelling in text generation, image processing, and diverse tasks including natural language understanding, programming, and problem-solving.

### A.4 ICL Prompt Selection Baselines

Here, we provide additional details on the two baselines used in our experiments.

**Random Sampling** This approach is a simple baseline where demonstration samples are randomly selected from the training set without considering any specific criteria or relevance to the target

example. Despite its simplicity and low computational cost, random sampling may not always select the most informative or relevant samples. By comparing our proposed method to random sampling, we can assess whether our approach significantly improves over a simple and unbiased selection strategy.

**BM25** This is a prominent word overlap-based method widely used in information retrieval tasks. It calculates the relevance score between a query and a document based on the frequency and importance of the query terms appearing. In the context of sample selection for ICL, BM25 selects demonstration samples with a high word overlap with the target example, aiming to select the most relevant samples based on their textual similarity. By comparing our proposed method to BM25, we can evaluate whether our AST-ICL approach provides additional benefits over a word overlap-based method in the context of SQL2Text tasks.

### A.5 Details of Experimental Cost

The costs of running the GPT-4 Experiments are quantified in Table 6. In addition, the total cost of running GPT-4o generation of Iterative ICL datasets was \$6.87 almost evenly split across the three datasets.

GPT-4 Experiments	
GPT-4 Zero-shot	1.51
GPT-4 + BM25-2	1.79
GPT-4 + AST-ICL-TOP-2	1.68
GPT-4 + Random-2	1.89
<hr/>	
GPT-4o Datasets Generation	6.86
<b>Total</b>	\$13.73

Table 6: Total USD Cost of GPT-4 Experiments

The prompt used for the ICL experiments run on GPT-4 is as follows:

<sup>5</sup>GPT-J Huggingface: EleutherAI/gpt-j-6b

<sup>6</sup>Mistral Huggingface: mistralai/Mistral-7B-v0.1

<sup>7</sup>CodeLlama Huggingface: meta-llama/CodeLlama-7b-hf




Figure 4: The total time (s) to generate the prompt samples for ICL by method, and dataset. Smaller datasets demonstrate a shorter time

#### GPT-4 Prompt

**System:** "You are an intelligent SQL Code assistant who effectively translates the intent and logic of the SQL queries into natural language that is easy to understand.  
**User:** Convert the given SQL query into a clear and concise natural language query limited to 1 sentence. Ensure that the request accurately represents the actions specified in the SQL query and is easy to understand for someone without technical knowledge of SQL.

## A.6 Experimental Environment

The `sqlglot`<sup>8</sup> Python library was used to parse SQL statements. We used Python 3.11.8 and PyTorch 2.2.1 with CUDA v12.4, and PyTorch Geometric 2.5.3. All experiments were run on (1) NVIDIA A100 32GB GPU.

## A.7 Implementation Details

**Tokenizer** A simple index tokenizer converts the SQL syntax tree node elements to a numeric form. Torch nn.Embedding was used to embed the token to a vector of 100. The edge index was constructed as directed edges as depicted in Figure 2.

Two GCN aggregation layers were used to aggregate the node embeddings followed by a global mean pool and a linear layer into vector outputs in a 2-dim space. For the AST-ICL variation, the vector outputs are clustered using k-means with 20 clusters. The number of clusters was configured based on the Silhouette score.

Dataset	Count
CoSQL-S2T	52
SparC-S2T	68
Spider-S2T	65

Table 7: The number of utterance improvements between the initial and the final generation using Iterative ICL.

## A.8 RQ1 Results

Table 7 outlines the number of utterances improved using Iterative ICL’s approach based on the generated and final refined variations.

Table 8 shows the effect on the measured scores when the newly generated datasets are used vs. using the datasets in their original form.

LLM	Dataset	AST-ICL-TOP-2 Results		
		BERTScore-1	BERTScore-2	BLEU
GPT-J	CoSQL-S2T	0.8153	0.7783	0.2813
	CoSQL	0.7584	0.7135	0.1914
Mistral	CoSQL-S2T	0.8632	0.8306	0.3732
	CoSQL	0.8194	0.7668	0.2417
CodeLlama	CoSQL-S2T	0.8571	0.8217	0.3578
	CoSQL	0.7933	0.7407	0.2012

Table 8: Test Results for AST-ICL-TOP-2 using original CoSQL and newly-introduced CoSQL-S2T dataset for SQL2Text

## A.9 RQ2 Results

Table 9 demonstrates the effect of the number of demonstration samples on the results. Figure 4 outlines the training time requirements for the various methods used.

## A.10 Iterative ICL Prompt

The Iterative ICL prompt used to generate the new datasets is outlined in Appendix A.11 below.

## A.11 Additional Experiments

Table 11 presents results based on GNN architecture. Our choice of GCN is motivated by its Align-Score results on which it has a slight edge (5 out of 9) compared to other options. As future work, it would be interesting to explore the other architectures comprehensively.

Table 10 presents the results based on the SQL query type. While nested and aggregate queries

<sup>8</sup><https://github.com/tobymao/sqlglot>

Samples Size	Method	GPT-J				Mistral				CodeLlama			
		BERTScore-1	BERTScore-2	BLEU	AlignScore	BERTScore-1	BERTScore-2	BLEU	AlignScore	BERTScore-1	BERTScore-2	BLEU	AlignScore
Two	Random-2	0.7837	0.7392	0.2021	0.4943	0.8515	0.8145	0.3509	0.7335	0.8384	0.8019	0.2996	0.6756
	BM25-2	0.7779	0.7421	0.2197	0.5189	0.8498	0.8143	0.3415	0.7499	0.835	0.8022	0.3314	0.7153
	<b>AST-ICL-TOP-2</b>	<b>0.8153</b> <sup>†‡</sup>	<b>0.7783</b> <sup>†‡</sup>	<b>0.2813</b> <sup>†‡</sup>	<b>0.6357</b> <sup>†‡</sup>	<b>0.8632</b>	<b>0.8306</b> <sup>†‡</sup>	<b>0.3732</b>	<b>0.8103</b> <sup>†‡</sup>	<b>0.8571</b> <sup>†‡</sup>	<b>0.8217</b> <sup>†‡</sup>	<b>0.3578</b> <sup>†</sup>	<b>0.7766</b> <sup>†‡</sup>
	<b>AST-ICL-2</b>	0.8082 <sup>†‡</sup>	0.7711 <sup>†‡</sup>	0.2376	0.5583	0.8577	0.82	0.3369	0.7382	0.8451	0.8085	0.3218	0.7031
Four	Random-4	0.8032	0.766	0.2424	0.5247	0.8663	0.8288	0.3646	0.7474	0.8534	0.8209	0.331	0.7185
	BM25-4	0.8101	0.7767	0.259	0.5671	0.8643	0.8298	0.3674	0.7767	0.8478	0.8177	<b>0.3566</b>	0.7402
	<b>AST-ICL-TOP-4</b>	<b>0.8333</b> <sup>†‡</sup>	<b>0.7974</b> <sup>†‡</sup>	<b>0.2964</b> <sup>†</sup>	<b>0.6457</b> <sup>†‡</sup>	<b>0.8722</b>	<b>0.8331</b>	<b>0.3799</b>	<b>0.7978</b>	<b>0.8667</b> <sup>‡</sup>	<b>0.832</b>	0.356	<b>0.7865</b> <sup>†</sup>
	<b>AST-ICL-4</b>	0.8171	0.7778	0.2601	0.5745	0.862	0.834	0.3851	0.7674	0.8557	0.8311	0.339	0.7395
Eight	Random-8	0.8176	0.7813	0.2462	0.5955	0.8717	0.8388	0.3952	0.7949	0.86	0.8278	0.3623	0.7481
	BM25-8	0.824	0.7944	0.3 <sup>†</sup>	0.5748	0.8785	0.8452	0.396	0.8362	0.8664	<b>0.8382</b>	<b>0.3935</b>	<b>0.7888</b>
	<b>AST-ICL-TOP-8</b>	<b>0.835</b> <sup>†</sup>	<b>0.7994</b> <sup>†</sup>	<b>0.3125</b> <sup>†</sup>	<b>0.6272</b>	0.8722	0.838	<b>0.3988</b>	0.7958	<b>0.8699</b>	0.8327	0.3709	0.7807
	<b>AST-ICL-8</b>	0.8294	0.7903	0.2785	0.6079	0.872	0.839	0.4112	0.7868	0.8658	0.8343	0.3611	0.7745

Table 9: Results against different demonstration sample sizes using the CoSQL-S2T dataset. The number after the hyphen refers to the number of demonstration samples used in the prompt, which is two in this experiment. <sup>†</sup> denotes t-test statistical significance compared to Random. <sup>‡</sup> denotes t-test statistical significance compared to BM25.

have lower zero-shot results, they benefit significantly from a higher samples size. In addition, AST-ICL provides the most improvement gains on these query types.

## Chain of Commentation Prompt

**System Prompt:** Given a user SQL query, follow these 3 steps to generate, review, and refine variations of utterances or questions that can be used to create the query accurately in a text-to-SQL generation task.

Step 1: Generate three utterances or question variations for the SQL query

- Produce three different but accurate variations of utterances or questions for the user's original SQL query.
- Ensure each variation presents a unique phrasing while maintaining the original query's intent, logic, key elements, and meaning.

Example:

```
{
  "Original SQL Query": "SELECT * FROM Flights WHERE city_from = 'New York' AND city_to = 'Boston' ORDER BY price ASC LIMIT 1;",
  "Generated Variations": [
    "What is the most affordable flight from New York to Boston?",
    "Can you show me the cheapest flight available between New York and Boston?",
    "Find the lowest-priced flight from New York to Boston."
  ]
}
```

In Step 2, review the generated variations by identifying any errors, ambiguities, or areas for improvement in each variation. Focus on accuracy, clarity, and alignment with the original SQL query's intent, and provide feedback for each variation.

Example:

```
{
  "Review Feedback": {
    "Variation 1": {
      "Feedback": "This variation is clear but could explicitly mention it is a search for a flight."
    },
    "Variation 2": {
      "Feedback": "Good phrasing, but it could specify 'currently available' for more precision."
    },
    "Variation 3": {
      "Feedback": "Accurate and concise. No changes needed."
    }
  }
}
```

In Step 3, refine and finalize the utterances by improving each variation based on the feedback in Step 2, ensuring better accuracy, clarity, fluency, and relevance to the original SQL query. Finalize the three refined versions of the utterances.

Example:

```
{
  "Final Refined Variations": [
    "What is the most affordable flight currently available from New York to Boston?",
    "Can you show me the cheapest flight available between New York and Boston?",
    "Find the lowest-priced flight from New York to Boston in the current listings."
  ]
}
```

Now, fill out the form below based on the Original SQL Query input from the user:

```
{
  "Original SQL Query": "<User's SQL Query Input>",
  "Generated Variations": [
    "<Generated Variation 1>",
    "<Generated Variation 2>",
    "<Generated Variation 3>"
  ]
}
```

For Step 2:

```
{
  "Review Feedback": {
    "Variation 1": {
      "Feedback": "<Feedback for Variation 1>"
    },
    "Variation 2": {
      "Feedback": "<Feedback for Variation 2>"
    },
    "Variation 3": {
      "Feedback": "<Feedback for Variation 3>"
    }
  }
}
```

For Step 3:

```
{
  "Final Refined Variations": [
    "<Refined Variation 1>",
    "<Refined Variation 2>",
    "<Refined Variation 3>"
  ]
}
```



Samples Size	Method	Simple (129 instances)				Nested (42 instances)				Aggregate (119 instances)			
		BERTScore-1	BERTScore-2	BLEU	AlignScore	BERTScore-1	BERTScore-2	BLEU	AlignScore	BERTScore-1	BERTScore-2	BLEU	AlignScore
Zero	Zero-Shot	0.5994	0.5125	0.0115	0.514	0.5718	0.4795	0.0065	0.2258	0.5693	0.4679	0	0.4804
Two	Random-2	0.766	0.7216	0.174	0.4306	0.7578	0.719	0.1645	0.4452	0.8119	0.7654	0.2458	0.5807
	BM25-2	0.7453	0.7141	0.1612	0.4285	0.8118 <sup>†</sup>	0.7791 <sup>†</sup>	0.2735 <sup>‡</sup>	0.5747	0.8013	0.7594	0.2642	0.5971
	AST-ICL-TOP-2	<b>0.804</b> <sup>†‡</sup>	<b>0.7799</b> <sup>†‡</sup>	<b>0.2463</b> <sup>†‡</sup>	<b>0.6052</b> <sup>†‡</sup>	<b>0.8438</b> <sup>†</sup>	<b>0.8107</b> <sup>†</sup>	<b>0.3572</b> <sup>†</sup>	<b>0.7186</b> <sup>†</sup>	<b>0.83</b>	<b>0.7981</b> <sup>†‡</sup>	<b>0.2959</b>	<b>0.702</b> <sup>†‡</sup>
	AST-ICL-2	0.7963 <sup>†‡</sup>	0.7536 <sup>†‡</sup>	0.2088	0.5428 <sup>†‡</sup>	0.8098 <sup>†</sup>	0.7828 <sup>†</sup>	0.2488	0.4666	0.8204	0.7859	0.2649	0.6076
Four	Random-4	0.7733	0.7465	0.1985	0.457	0.8029	0.7596	0.2381	0.5299	0.8357	0.7893	0.2916	0.5963
	BM25-4	0.7916	0.7665	0.2189	0.5252	0.8274	0.7996 <sup>†</sup>	0.272	0.5711	0.8241	0.7796	0.298	0.6111
	AST-ICL-TOP-4	<b>0.8145</b> <sup>†</sup>	<b>0.783</b> <sup>†</sup>	<b>0.2483</b>	<b>0.575</b> <sup>†</sup>	<b>0.8435</b> <sup>†</sup>	<b>0.8039</b> <sup>†</sup>	<b>0.3038</b>	<b>0.6529</b>	<b>0.8501</b> <sup>†</sup>	<b>0.8107</b> <sup>†</sup>	<b>0.3461</b>	<b>0.7197</b> <sup>†‡</sup>
	AST-ICL-4	0.7959	0.7606	0.2085	0.4918	0.828	0.8044 <sup>†</sup>	0.2984	0.6155	0.8362	0.7871	0.3026	0.6496
Eight	Random-8	0.7963	0.7669	0.2234	0.5451	0.8213	0.7892	0.2782	0.5667	0.8395	0.7941	0.2597	0.6603
	BM25-8	0.8054	0.7788	0.2521	0.5057	<b>0.8513</b>	<b>0.8273</b>	<b>0.3627</b>	0.6204	0.8345	<b>0.7998</b>	<b>0.3298</b>	0.6337
	AST-ICL-TOP-8	<b>0.8246</b> <sup>†</sup>	<b>0.7971</b> <sup>†</sup>	<b>0.3007</b> <sup>†</sup>	<b>0.5693</b>	0.8463	0.8175	0.327	<b>0.6227</b>	<b>0.8423</b>	0.7956	0.3201	<b>0.6915</b>
	AST-ICL-8	0.8104	0.7774	0.2501	0.5403	0.838	0.8048	0.3133	0.5784	0.847	0.7991	0.2971	0.6915

Table 10: Results based on query type and number of demonstrations on GPT-J using the CoSQL-S2T dataset. The number after the hyphen refers to the number of demonstration samples used in the prompt. <sup>†</sup> denotes t-test statistical significance compared to Random. <sup>‡</sup> denotes t-test statistical significance compared to BM25.

Samples Size	Graph Method	GPT-J				Mistral				CodeLlama			
		BERTScore-1	BERTScore-2	BLEU	AlignScore	BERTScore-1	BERTScore-2	BLEU	AlignScore	BERTScore-1	BERTScore-2	BLEU	AlignScore
Two	GCN	0.8153	0.7783	0.2813	0.6357	0.8632	0.8306	0.3732	0.8103	0.8571	0.8217	0.3578	0.7766
	GraphSAGE	0.8239	0.7896	0.282	0.6307	0.8711	0.8409	0.3929	0.8074	0.8671	0.8376	0.3661	0.8033
	GAT	0.8204	0.7918	0.2827	0.6614	0.8646	0.8313	0.372	0.7994	0.8536	0.8217	0.3557	0.7639
Four	GCN	0.8333	0.7974	0.2964	0.6457	0.8722	0.8331	0.3799	0.7978	0.8667	0.832	0.356	0.7865
	GraphSAGE	0.8373	0.8001	0.3073	0.6284	0.8768	0.8468	0.3966	0.8242	0.8706	0.8384	0.3768	0.7953
	GAT	0.8258	0.792	0.3072	0.6103	0.8657	0.831	0.3851	0.7924	0.8621	0.834	0.3652	0.7762
Eight	GCN	0.835	0.7994	0.3125	0.6272	0.8722	0.838	0.3988	0.7958	0.8699	0.8327	0.3709	0.7807
	GraphSAGE	0.8444	0.8107	0.3114	0.6306	0.8817	0.8523	0.4203	0.8209	0.8712	0.8427	0.3818	0.784
	GAT	0.8339	0.8007	0.2979	0.6118	0.8741	0.8425	0.404	0.8156	0.871	0.8415	0.378	0.8183

Table 11: Results based on GNN type and number of demonstrations on using the CoSQL-S2T dataset. <sup>†</sup> denotes t-test statistical significance compared to Random. <sup>‡</sup> denotes t-test statistical significance compared to BM25.