

# KU-DMIS at EHRSQL 2024: Generating SQL query via question templatization in EHR

Hajung Kim<sup>1\*</sup>, Chanhwi Kim<sup>1\*</sup>, Hoonick Lee<sup>1</sup>, Kyochul Jang<sup>1</sup>, Jiwoo Lee<sup>1</sup>,  
Kyungjae Lee<sup>2</sup>, Gangwoo Kim<sup>1</sup>, Jaewoo Kang<sup>1,3†</sup>

<sup>1</sup>Korea University, <sup>2</sup>LG AI Research, <sup>3</sup>AIGEN Sciences  
{hajungk, chanhwi\_kim, hoonick, gcj0125, hijiwoo7}@korea.ac.kr  
kyungjae.lee@lgresearch.ai, {gangwoo\_kim, kangj}@korea.ac.kr

## Abstract

Transforming natural language questions into SQL queries is crucial for precise data retrieval from electronic health record (EHR) databases. A significant challenge in this process is detecting and rejecting unanswerable questions that request information beyond the database’s scope or exceed the system’s capabilities. In this paper, we introduce a novel text-to-SQL framework that robustly handles out-of-domain questions and verifies the generated queries with query execution. Our framework begins by standardizing the structure of questions into a templated format. We use a powerful large language model (LLM), fine-tuned GPT-3.5 with detailed prompts involving the table schemas of the EHR database system. Our experimental results demonstrate the effectiveness of our framework on the EHRSQL-2024 benchmark benchmark, a shared task in the ClinicalNLP workshop. Although a straightforward fine-tuning of GPT shows promising results on the development set, it struggled with the out-of-domain questions in the test set. With our framework, we improve our system’s adaptability and achieve competitive performances in the official leaderboard of the EHRSQL-2024 challenge.

## 1 Introduction

Electronic Health Records (EHRs) are crucial elements of the contemporary healthcare system, storing patients’ medical histories in relational databases. However, retrieving data from EHRs can be challenging, requiring specialized training in Structured Query Language (SQL). To bridge this gap, previous works build AI-powered systems that parse the user’s question (Yin et al., 2020; Brunner and Stockinger, 2021) or convert it into an SQL query that the database can process. Lee et al. (2022) identify an essential component in this

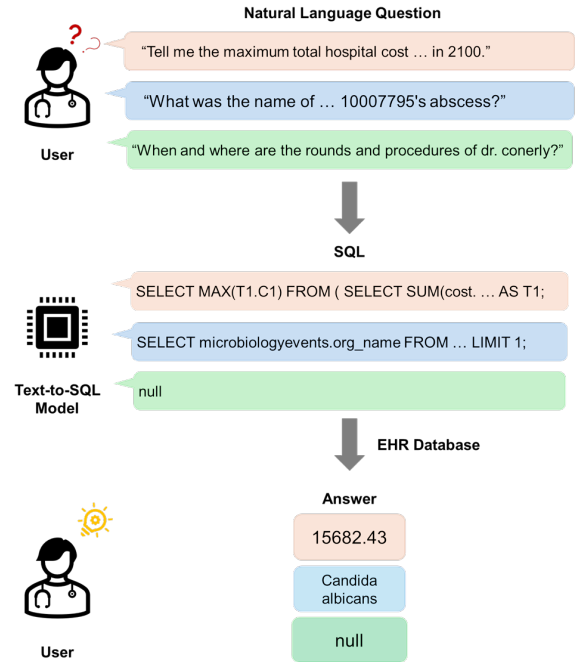


Figure 1: In the proposed Text-to-SQL framework, when a query is presented in natural language, the model generates SQL code to retrieve the required information from the database. If the query requires information absent from the database, the Text-to-SQL model returns a 'null' response.

text-to-SQL task; recognizing and adequately handling unanswerable questions that seek information beyond what the database contains. Hence, to ensure reliability and trustworthiness, systems should be able to refrain from answering unanswerable questions.

To further encourage research in this field, the Clinical NLP 2024 workshop has introduced a new shared task called EHRSQL-2024 (Lee et al., 2024) to motivate the development of more reliable question-answering (QA) systems. The EHRSQL-2024 dataset involves the real-world needs of medical personnel, incorporating templates of their most common questions. In this challenge, systems are

\* Equal contribution, † Corresponding author

tasked to generate SQL queries that accurately return the desired information from tables from the MIMIC-IV (Johnson et al., 2016). Additionally, the dataset includes inherently unanswerable questions, either due to the restrictions of the database schema or the request for information not contained within the databases. On the other hand, the test set presents distracting question types that contain noisy words, further testing the robustness of participants’ systems.

In this paper, we introduce a novel framework created to convert natural language questions to SQL queries for EHR databases. This framework transforms free-form questions into a templated format to handle distracting questions. We fine-tune GPT-3.5-turbo (Brown et al., 2020), one of the most performant large language models (LLMs), optimizing it to effectively interpret intricate medical queries and produce the corresponding SQL queries. We also provide detailed prompts that describe the tables in the EHR database.

For SQL generation, given the task’s complexity and relationships between tables, we break it down into two steps: selecting relevant tables and then generating SQL by reflecting in-depth on the selected tables. We enhance the accuracy and reliability of the generated SQL queries by correcting any errors in table names and applying ensemble techniques with majority voting.

Our empirical results of fine-tuned GPT-3.5 on the EHRSQL-2024 benchmark highlight its capability, achieving third place on the development set. However, it revealed a limitation in generalizing to questions in the test set that diverged from the predefined templates. By using our framework, we successfully address this gap between free-form questions, resulting in a notable improvement of 26.5 in the RS(10) metric in the test set. Additionally, we find that decomposing the task into two steps contributed to this success, with a significant improvement in RS (10) in the test set. Furthermore, by employing further verification and ensemble techniques, we attained fourth place in the EHRSQL-2024 challenge’s official leaderboard.

We conduct in-depth analyses of the questions to uncover disparities in each split. In particular, we apply  $N$ -gram counting of the questions to highlight the distribution gaps. This variation emphasizes the need to develop a resilient model capable of adapting to and performing consistently across datasets with diverse word distributions. Additionally, we manually categorize the unanswer-

able questions into three distinct types.

## 2 Related Works

### 2.1 Text-to-SQL Generation

Text-to-SQL conversion requires interpreting natural language questions, matching them with the database schema, and producing accurate SQL queries that reflect the question’s intent. This task is particularly challenging for individuals unfamiliar with database structures, highlighting the need for methods that translate natural language into SQL queries—a focus of ongoing research due to real-world applications. However, accurately generating SQL code from natural language is complex, mainly because of the challenges in integrating precise database knowledge into the model (Qin et al., 2022; Katsogiannis-Meimarakis and Koutrika, 2023).

Initially, efforts to address Text-to-SQL employed predefined rules (Sen et al., 2020) to handle existing difficulties. The field has evolved since then to explore encoder-decoder models (Cai et al., 2017; Popescu et al., 2022), and Text-to-SQL is tested on sequence-to-sequence approaches (Qi et al., 2022). With the rapid advancement in deep learning research, methodologies incorporating graph representation (Xu et al., 2018; Wang et al., 2019; Brock et al., 2021) and attention mechanisms (Liu et al., 2023b) have been extensively applied to Text-to-SQL tasks. Additionally, the Text-to-SQL task, tailored to real-world data, has been conducted on datasets such as WikiSQL (Zhong et al., 2017), Spider (Yu et al., 2018), KaggleD-BQA (Lee et al., 2021), and BIRD (Li et al., 2023).

With the emergence of LLMs like GPT (Brown et al., 2020) and LLaMA (Touvron et al., 2023), research leveraging these models has proliferated. Their comprehensive pretraining on massive text corpora enables them to show promising results using techniques like prompt engineering and in-context learning (Trummer, 2022; Liu et al., 2023a; Chang and Fosler-Lussier, 2023; Dong et al., 2023; Sun et al., 2023). Despite these advancements, exploring supervised fine-tuning has led to even greater enhancements in their performance (Gao et al., 2023).

### 2.2 Text-to-SQL in EHR database

The MIMIC-III (Johnson et al., 2016) is a prominent EHR database in the healthcare domain. MIMICSQL (Tarbell et al., 2023) is the first dataset

constructed based on the MIMIC-III database, designing questions generated from pre-formatted templates. Similarly, emrKBQA (Raghavan et al., 2021) derived from the MIMIC-III database and the emrQA (Yue et al., 2020) dataset focused on clinical reading comprehension expands the research scope. EHRSQL, introduced by Lee et al. (2022), is an extensive text-to-SQL dataset that is linked to the two open-source EHR databases, MIMIC-III and eICU (Pollard et al., 2018). Created based on feedback from 222 professionals with varied experience levels, EHRSQL covers a wide range of real-world scenarios. This dataset includes time-sensitive questions to highlight the critical importance of time in the healthcare domain. Additionally, it incorporates unanswerable questions to evaluate the system’s capability to recognize and handle such inquiries effectively.

### 2.3 Discriminating Unanswerable Questions

The distinction between answerable and unanswerable questions is crucial in NLP tasks, especially in domains where accuracy and reliability are critical, such as healthcare. Discriminating between these types of questions is complex due to the subtle differences in what a question may require for a satisfactory answer. The language models often exhibit overconfidence in their ability to accurately respond to a given question. To address this, the specialized datasets have been designed through various methodologies, such as rule-based editing (Jia and Liang, 2017), distant supervision (Joshi et al., 2017), and crowdsourcing (Rajpurkar et al., 2018), each method offering its own set of benefits and challenges for identifying unanswerable questions. This advancement facilitates more reliable and accurate question-answering capabilities, which is crucial for applications where the cost of misinformation can be high.

## 3 Methods

Figure 2 presents an outline of our proposed methodology. Our process starts with the templating of questions, transforming free-form inquiries into a standardized format to ensure consistency in how queries are represented. Additionally, we enrich the model’s understanding by supplying detailed information about the database tables, thereby improving its capacity to formulate precise queries. To further elevate the accuracy of the generated SQL queries, we introduce a verification

phase to confirm that the queries accurately correspond to the intended data retrieval objectives.

### 3.1 Question Templating

We introduce question templating to handle the diverse forms of question presentation. This approach addresses the challenge of questions deviating from a standard template by employing a reverse engineering strategy. By converting free-form questions into a templated format, we aim to align them more closely with similar patterns, thus bridging the gap between the varied question formats in real-world contexts. Specifically, we prompt GPT-4-turbo to rewrite questions to match the structure of pre-defined templates more closely.

Identifying semantically close questions involves searching for questions similar to the input question. This similarity is quantified by calculating the Euclidean distance between the question embeddings and comparing input questions to potential neighbors. We mask identification information to ensure that specific table values do not skew this comparison. For example, a question like *"Count how many times in the first hospital visit patient 10004457 had coronary arteriography using two catheters."* is transformed into *"Count how many times patient <patient number> had <procedure> during their first hospital visit."* By adopting this method, we achieve a uniform question format, effectively standardizing free-form queries and reducing discrepancies in dataset distribution. The templated question is utilized as the input question.

### 3.2 SQL generation

Considering the complexity of the text-to-SQL task and the intricate relationships among more than ten tables in the database, we propose a two-stage approach that involves a table selection phase followed by a self-reflection phase.

**Table Selection** We task the GPT model with converting natural language questions into SQL queries. The construction of prompts for the model involves three essential components: (1) outlining the text-to-SQL task guides the model to convert a natural language question into an SQL query for data retrieval from the EHR database. We clarify that the database uses SQLite and highlight the syn-

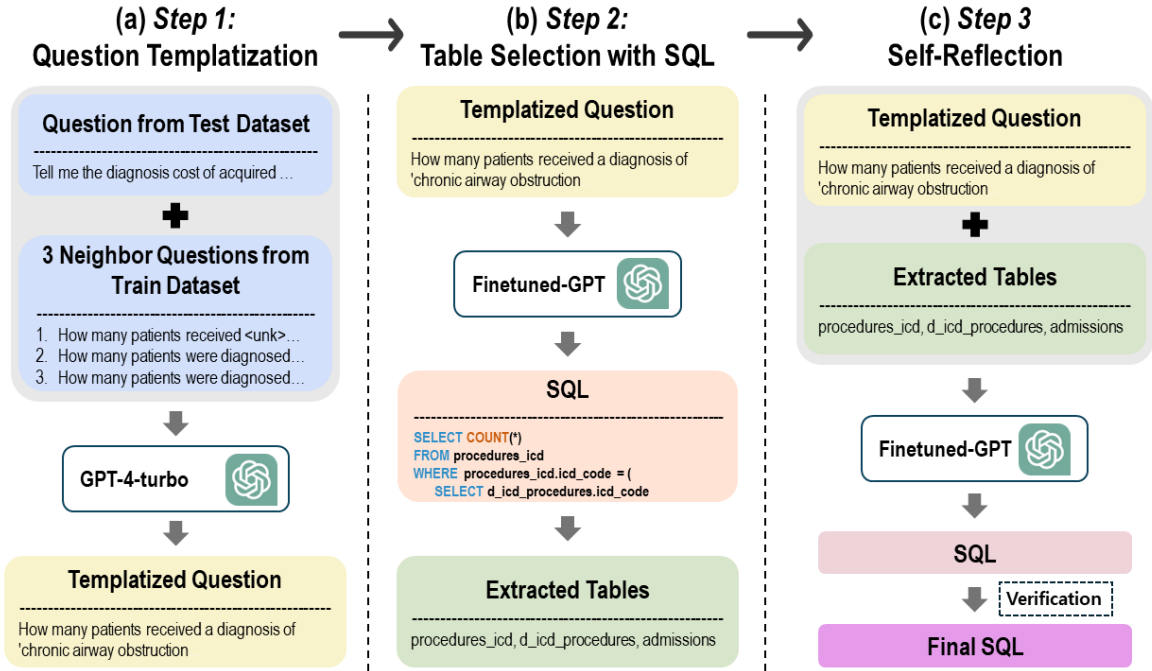


Figure 2: Overview of our framework. (a) Question Templatzation (Sec. 3.1). Implementing question templatzation to convert free-form questions into a structured format. (b) SQL Generation (Sec. 3.2). Providing task outlines and table information to aid in precise query generation. (c) Self-Reflection and Verification (Sec. 3.2, 3.3). Providing detailed table information identified in the initial SQL generation and then finalizing the process.

tactical nuances between SQLite and other SQL dialects to guide the model’s syntax choice. (2) By detailing the database tables, we describe the database’s complex structure, listing over ten tables with brief descriptions and their respective columns. This detail is crucial since it aids the model in identifying the relevant tables and navigating their relational schema without direct access to the database values. We follow the format introduced in DAIL-SQL (Gao et al., 2023) for table schema details, which allows both natural language and SQL representations. (3) Presenting the question for conversion is the natural language question to be transformed. By using this prompt, we prompt the model to produce an SQL query that matches the question and subsequently identifies the table name mentioned within the SQL query.

**Self-Reflection** The prompt for the self-reflection stage is similar to Table Selection, except for detailing table information. In this stage, the prompt is augmented with detailed descriptions for each table column identified in the initial SQL query. This refinement aims to enhance the SQL query formulation by providing a more comprehensive understanding of the selected table’s specifics, enabling the model to generate a

more accurate and targeted SQL query.

### 3.3 SQL Verification

We implement a verification step on the generated SQL queries to address two specific scenarios. In the first scenario, some questions can be technically converted into SQL queries but remain unanswerable due to the absence of required information in the dataset. To avoid providing incorrect SQL results and improper answers, which are unanswerable, we verify the validity of each SQL query by executing it against the database. If the execution results in an error, indicating the SQL query cannot retrieve the correct answer, we replace the SQL query with *null* instead. This adjustment ensures the query is considered valid but unanswerable, optimizing score outcomes.

The second scenario addresses instances where the generated SQL query includes incorrect references to table names or column names. In such cases, we identify the correct table name and associated column names based on the table values mentioned in the SQL query. We then modify the SQL query to accurately reflect the proper table name and column names to which the table values correspond. This correction process

Team	Development				Test			
	RS(0)	RS(5)	RS(10)	RS(N)	RS(0)	RS(5)	RS(10)	RS(N)
LG AI Research & KAIST	90.37	<b>89.51</b>	<b>88.65</b>	<b>-109.6</b>	<b>88.17</b>	<b>84.75</b>	<b>81.32</b>	<b>-711.83</b>
PromptMind	66.38	59.5	52.62	-1533.62	82.6	78.75	74.89	-817.4
ProbGate	84.18	79.45	74.72	-1015.82	81.92	78.06	74.21	-818.08
<b>KU-DMIS (Ours)</b>	<b>91.57</b>	82.98	74.38	-1908.43	72.07	65.64	59.21	-1427.93
oleg1996	47.03	34.14	21.24	-2952.97	68.89	56.47	44.04	-2831.11
LTRC-IIITH	N/A	N/A	N/A	N/A	66.84	55.27	43.7	-2633.16
Saama Technologies	57.78	50.47	43.16	-1642.22	53.21	44.64	36.08	-1946.79
TEAM_optimist	N/A	N/A	N/A	N/A	14.14	-349.61	-713.37	-84 885.86

Table 1: Official results of the leaderboard on EHRSQL-2024 dataset. The teams are ranked based on Reliability Score RS(10).

ensures that the SQL query accurately represents the intended data retrieval operation, aligning with the database’s schema.

### 3.4 Ensemble with Majority Voting

We incorporate an ensemble method to determine the final SQL query. We first instruct GPT-4-turbo to evaluate whether the generated SQL query accurately captures the intent of the original natural language question. This alignment check ensures that the model prioritizes the core intent of the query, such as using the ‘COUNT’ function in SQL queries asking for a count of patients. To finalize the SQL query or its resulting answer from the database execution, we adopt a majority voting system. This ensemble strategy mitigates the variability inherent in the fine-tuned model and improves the robustness. Using majority voting to select the SQL query or derive its answer aims to improve performance metrics by effectively managing *null* responses.

## 4 Experiments

### 4.1 Experimental Setup

**Dataset** We evaluate our framework using the EHRSQL-2024 challenge benchmark dataset (Lee et al., 2024). This large-scale Text-to-SQL dataset contains 5,124 instances in the train set, 1,163 instances in the development set, and 1,167 instances in the test set, spanning 17 tables. The train dataset comprises natural language questions paired with their corresponding SQL queries. However, the SQL queries associated with the questions in the development and test sets are not provided.

**Metric** Following Lee et al. (2024), we use the Reliability Score (RS). RS is unique because it rewards correct SQL queries for answerable questions ( $Q_{ans}$ ) and the choice to abstain from answering unanswerable questions ( $Q_{una}$ ). At the same time, it penalizes incorrect SQL generation for  $Q_{ans}$  and any attempt to answer  $Q_{una}$ . Moreover, RS includes a penalty factor ‘c’ to adjust the evaluation’s strictness according to specific safety requirements. The corresponding formula is as follows.

$$\phi_c(x) = \begin{cases} 1 & \text{if } x \in Q_{ans}, g(x) = 1, \text{Acc}(x) = 1 \\ 0 & \text{if } x \in Q_{ans}, g(x) = 0 \\ -c & \text{if } x \in Q_{ans}, g(x) = 1, \text{Acc}(x) = 0 \\ -c & \text{if } x \in Q_{una}, g(x) = 1 \\ 1 & \text{if } x \in Q_{una}, g(x) = 0 \end{cases}$$

The adaptability of RS is demonstrated by evaluating models under four different scenarios, which vary based on the severity of the penalty term: RS(0), RS(10), and RS(N). For this challenge, the primary metric is RS(10), emphasizing the importance of accurately assessing answerable questions and refraining from generating SQL for unanswerable questions.

### 4.2 Implementation Details

We utilize GPT, one of the most performant Large Language Models (LLMs), to enhance the translation from text to SQL. We investigate the effectiveness of in-context learning and supervised fine-tuning methods. We employ GPT-3.5-turbo, GPT-4-turbo, and GPT-4 models for in-context learning, augmenting the prompt with three more examples. These examples consist of pairs of

Model	Few-shot	Table info.	RS(0)	RS(10)	RS(N)
GPT-3.5-turbo	0	O	29.53	-250.19	-28670.47
	3	O	48.54	15.40	-3351.46
	3	-	75.34	29.53	-4624.66
GPT-4-turbo	0	O	36.74	-196.20	-23863.26
	3	O	67.84	-18.91	-8832.16
	3	-	85.87	42.98	-4314.13
GPT-4	0	O	38.40	-215.98	-26061.60
	3	O	79.82	15.50	-6520.18
	3	-	90.45	62.18	-2809.55
Finetuned-GPT	-	O	<b>98.05</b>	<b>91.23</b>	<b>-601.95</b>

Table 2: Training set performance. Comparison of GPT models.

semantically related questions, including the input question and their corresponding SQL queries. The semantic relatedness is determined by calculating the Euclidean distance between question embeddings derived from the training dataset and the input question embedding. For supervised fine-tuning, we focus on the GPT-3.5-turbo model, the primary model available for fine-tuning. The model is prompted without including neighboring examples. Based on the evaluation results, it is clear that the supervised fine-tuning methodology is particularly effective in addressing the challenges inherent in text-to-SQL tasks. Further details are provided in section 4.4. Consequently, the fine-tuned GPT-3.5-turbo model is selected for further detailed experiments.

### 4.3 Leaderboard Results

Table 1 presents the scores of the participants’ systems, ranked according to the RS(10) score. We secured fourth place in the test set rankings. All participating teams utilized Large Language Models (LLMs), with the top four teams, including ours, primarily employing a fine-tuned GPT model and incorporating various other techniques. This table underscores the efficacy of LLMs in addressing Text-to-SQL tasks.

### 4.4 In-Context Learning and Fine-tuning

To evaluate the effectiveness of various GPT models for Text-to-SQL tasks, we conduct experiments with GPT-3.5-turbo, GPT-4-turbo, and GPT-4 for in-context learning and a fine-tuned version of GPT-3.5-turbo for supervised fine-tuning. Due to submission limitations, we assessed the GPT models using the training set. We adopt a k-fold cross-validation method with  $k = 5$ , training on four

Model	RS(0)	RS(10)	RS(N)
GPT-3.5-turbo	70.34	13.59	-6529.66
GPT-4-turbo	76.53	-6.02	-9523.47
GPT-4	79.28	-17.88	-11220.72
Finetuned-GPT	<b>93.12</b>	50.99	-4806.88
w/ table info. in SQL form	83.23	17.02	-7616.77
w/ Self-Reflection	83.15	62.51	-2316.85
w/ Ensemble	91.57	<b>74.38</b>	<b>-1908.43</b>

Table 3: Ablation study conducted on the development set showcases the performance of in-context learning with few examples using GPT-3.5-turbo, GPT-4-turbo, and GPT-4, alongside fine-tuning performed with GPT-3.5-turbo using various additional techniques.

fold and evaluating the remaining fold. To maintain the balance of answerable and unanswerable questions in the training dataset, we divide unanswerable questions into three categories. When partitioning the training dataset into five folds, we ensured that the proportions of these categories were reflected in each fold. A detailed analysis of these categorized groups can be found in section 5.2.

Table 2 presents the comparison results of the GPT models. We experimented with variations by providing few-shot examples and including table information. The fine-tuned GPT model demonstrates superior performance across all metrics, making it our model of choice. Interestingly, the inclusion of table information slightly reduces performance in all in-context learning scenarios. We speculate that the table information in our experiment, which merely lists table names and column names, lacks detailed relational data like primary and foreign keys. Consequently, this minimal and potentially uninformative text might have acted as a distraction.

### 4.5 Table Information Format

The prompt includes table information to accurately identify the table and column names. Following the DAIL-SQL approach (Gao et al., 2023), we explore different formats of presenting table information, in both natural language and SQL format, within the same prompt framework. Our experiments, detailed in table 3, reveal that presenting table information in SQL format results in a decrease in the RS (10) score from 50.99 to 17.02. This suggests that natural language formats are more readily interpretable by the language model such as GPT.

Model	Inclusion	Jaccard	Exact Match
GPT-3.5-turbo	0.7933	0.7930	0.7836
GPT-4-turbo	0.8912	0.8908	0.8723
GPT-4	0.9250	0.9244	0.9123
<b>Finetuned-GPT</b>	<b>0.9857</b>	<b>0.9855</b>	<b>0.9844</b>
Table Selector (GPT-3.5-turbo)	0.8976	0.8488	0.7115

Table 4: Table selection performance.

Model	RS(0)	RS(10)	RS(N)
Finetuned-GPT (Ensemble)	72.07	<b>59.21</b>	<b>-1427.93</b>
Finetuned-GPT (Single)	78.06	13.80	-7421.94
w/ Self-Reflection	77.55	27.85	-5722.45
w/ Question Templatzation	<b>80.55</b>	40.27	-4619.45

Table 5: Ablation study on the test set. We provide the performance of ensembled and single results. Every component, including SQL regeneration and question templatzation, plays a key role in enhancing overall performance.

#### 4.6 Table Selection Results

Considering the complexity of over ten tables and the resulting SQL queries that reference multiple tables, we hypothesize that a self-reflection incorporating selective, detailed table information could enhance the accuracy of the generated SQL queries. In preparation for this self-reflection process, we assess the accuracy of the tables retrieved in the generated SQL queries. This assessment involves calculating the accuracy between correct tables and the tables extracted from the generated SQL queries. We use three metrics as an accuracy score: 1) inclusion score (indicating the presence of the correct tables within the generated SQL), 2) the Jaccard similarity score (comparing the intersection to the union of correct and extracted tables), 3) and the exact match score.

Table 4 suggests that the fine-tuned GPT model effectively identifies the relevant tables without a dedicated table selection model. We extract tables from the initially generated SQL queries and use prompts augmented with detailed information, such as descriptions of each column and examples of values, for the fine-tuned GPT model. The comparison between the initially generated SQL and the outcomes after the self-reflection stage, table 3 shows an increase in the RS(10) score from 50.99 to 62.51 in the development set, and table 5 also illustrates an improvement in the RS(10) score from 13.80 to 27.85 in the test set. This improvement indicates that the regenerated SQL queries provide more reliable and accurate outputs.

#### 4.7 Question Templatzation

Our analysis focuses on the characteristics of the questions across each dataset. It reveals a decline in the fine-tuned GPT model’s scores from the training set to the development and test sets. This pattern highlights substantial variations among the training, development, and test datasets. To mitigate these discrepancies, we employ the technique that reverses the deviation of questions from templates. We utilize GPT-4-turbo to rephrase the original question. By prompting GPT-4-turbo with the original question and semantically similar questions from the training set and template from (Lee et al., 2022), we aim to achieve consistency with related queries. This approach significantly reduces the distribution gaps between the training and test sets, as demonstrated in Table 5. The improvement in the RS(10) score from 13.80 to 40.27 highlights the effectiveness of question templatzation by comparing the performance of a single model before and after its application.

### 5 Analysis

In this section, we analyze the word distribution of questions for each dataset split: training, development, and test sets. The objective is to identify variations in question composition among these datasets. Furthermore, we investigate the distribution of unanswerable questions in the training set to better understand questions that yield an *null* response.

To focus solely on word analysis and minimize noise, we eliminate punctuation marks such as ".", ",", and "?", remove stop words such as "the", "a", and "an" from the questions, and convert all letters to lowercase. After eliminating these elements, we analyze the processed questions using N-grams. The analysis is limited to 1 to 3-grams, which is sufficient for understanding the context of questions while excluding the aforementioned noise. Appendix D details the ten most prevalent words alongside their respective frequencies within each dataset arranged in non-increasing order, including the collection of unanswerable queries labeled as *Unanswerable Train set*.

#### 5.1 N-gram Distribution

The initial three columns of Appendix D enumerates the top ten most frequent words in each dataset alongside their respective frequencies. Analysis of Appendix D indicated that words with high fre-

quency within one dataset tended to be frequent across other datasets as well, suggesting a pattern of similarity. However, it was observed that words with lower frequency, which were not included in the table, often did not appear in other datasets. This discrepancy became particularly evident within the context of 3-gram sets, highlighting a distinct distribution among the datasets.

This disparity underscores the necessity of developing a robust model that can adapt and excel across datasets with different word distributions.

## 5.2 Category of Unanswerable Questions

We analyzed the training set’s *null* distribution, identifying 450 unanswerable questions. Our initial qualitative analysis involved categorizing these *null*-labeled questions into three distinct groups through a detailed manual review: (1) Incorrect Patient Number, (2) Require External Knowledge, (3) Out of EHR Database.

In the first case, based on the MIMIC-IV dataset’s criteria, a legitimate patient number is identified by its 8-digit configuration; thus, questions featuring a patient number with fewer or more than 8 digits invariably resulted in a *null* response. Regarding the second case, specific questions, for example, *"I am curious what the protocols for the drugs that work to treat cancer."* could potentially be answered by a knowledgeable individual or through QA tasks using external information resources. The third group, while seemingly akin to the second, differed in that the questions could technically be converted into SQL queries; however, they remained unanswerable due to the absence of the required information in the dataset. Example questions include: *"Has patient 23224 an appointment in another hospital department?"*. Further examples for each category, along with their respective frequencies, are detailed in [Appendix A](#).

Additionally, a quantitative analysis of unanswerable questions was also conducted using N-grams. By examining the differences in word distribution between "answerable" and "unanswerable" questions, as highlighted by the contrast between the first and last columns of [Appendix D](#), significant disparities were noted. For instance, an examination of the 1-gram columns for both the training set and the Unanswerable Training set reveals that the only overlapping words are "patient" and "last." This indicates a significantly different distribution between the two datasets.

Based on both qualitative and quantitative analy-

sis, we were able to refine our framework to avoid generating SQL queries for questions that solely comprise words found in the unanswerable questions of the training set.

## 6 Conclusion

Throughout the challenge, we noticed that differences in the way data is distributed across training, development, and test sets can make it hard for our model to determine which questions are answerable or not. To tackle this issue, we templated questions to make the word distribution of development and test data more similar to the training data. This method aimed to bridge the gap between the datasets, helping the model better understand the features of unanswerable questions within the test dataset.

Although we did not address this in this study, we anticipate that future research could see performance improvements by augmenting the training dataset to more closely match the distribution of unanswerable questions in the development and test sets. Focusing on refining the test data to align more closely with the characteristics observed in the training datasets, we expect to increase model performance in identifying unanswerable questions. Such data augmentation strategies could bridge the remaining gaps between datasets and ensure a more robust model performance across varied datasets. Also, we utilized finetuned gpt-3.5-turbo, which is expensive and unusable for other researchers. Thus, further study should be done with open sourced models, like llama or gemma.

## Acknowledgments

This work was supported in part by the National Research Foundation of Korea [NRF2023R1A2C3004176], the Ministry of Health & Welfare, Republic of Korea [HR20C0021(3)], the Ministry of Science and ICT (MSIT) [RS-2023-00262002].

## References

Andrew Brock, J. Donahue, Karen Simonyan, Ruisheng Cao, Lu Chen, Zhi Chen, Yanbin Zhao, Donghyun Choi, M. Shin, EungGyun Kim, Xiang Deng, Ahmed Hassan Awadallah, Oleksandr Meek, Huan Polozov, Sun Matthew, Yujian Gan, Xinyun Chen, Qiuping Huang, John R Purver, Jinxia Woodward, Xie Peng-301, Amol Kelkar, Rohan Relan, V. Bhardwaj, and Saurabh Vaichal. 2021. [S 2 sql: Injecting syntax to question-schema interaction graph encoder](#)



- for text-to-sql parsers. In *Findings of the Association for Computational Linguistics: ACL 2022*.
- 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.
- Ursin Brunner and Kurt Stockinger. 2021. Valuenet: A natural language-to-sql system that learns from database information. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, pages 2177–2182. IEEE.
- Ruichu Cai, Boyan Xu, Zhenjie Zhang, Xiaoyan Yang, Zijian Li, and Zhihao Liang. 2017. An encoder-decoder framework translating natural language to database queries. In *International Joint Conference on Artificial Intelligence*.
- Shuaichen Chang and Eric Fosler-Lussier. 2023. How to prompt llms for text-to-sql: A study in zero-shot, single-domain, and cross-domain settings. *ArXiv*, abs/2305.11853.
- Xuemei Dong, C. Zhang, Yuhang Ge, Yuren Mao, Yunjun Gao, Lu Chen, Jinshu Lin, and Dongfang Lou. 2023. C3: Zero-shot text-to-sql with chatgpt. *ArXiv*, abs/2307.07306.
- Dawei Gao, Haibin Wang, Yaliang Li, Xiuyu Sun, Yichen Qian, Bolin Ding, and Jingren Zhou. 2023. Text-to-sql empowered by large language models: A benchmark evaluation. *ArXiv*, abs/2308.15363.
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. *ArXiv*, abs/1707.07328.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *ArXiv*, abs/1705.03551.
- George Katsogiannis-Meimarakis and Georgia Koutrika. 2023. A survey on deep learning approaches for text-to-sql. *The VLDB Journal*, 32:905–936.
- Chia-Hsuan Lee, Oleksandr Polozov, and Matthew Richardson. 2021. Kaggledbqa: Realistic evaluation of text-to-sql parsers. In *Annual Meeting of the Association for Computational Linguistics*.
- Gyubok Lee, Hyeonji Hwang, Seongsu Bae, Yeonsu Kwon, Woncheol Shin, Seongjun Yang, Minjoon Seo, Jong-Yeup Kim, and Edward Choi. 2022. Ehrsql: A practical text-to-sql benchmark for electronic health records. *Advances in Neural Information Processing Systems*, 35:15589–15601.
- Gyubok Lee, Sunjun Kweon, Seongsu Bae, and Edward Choi. 2024. Overview of the ehrsql 2024 shared task on reliable text-to-sql modeling on electronic health records. In *Proceedings of the 6th Clinical Natural Language Processing Workshop*, Mexico City, Mexico. Association for Computational Linguistics.
- Jinyang Li, Binyuan Hui, Ge Qu, Binhua Li, Jiayi Yang, Bowen Li, Bailin Wang, Bowen Qin, Rongyu Cao, Ruiying Geng, Nan Huo, Chenhao Ma, Kevin C. Chang, Fei Huang, Reynold Cheng, and Yongbin Li. 2023. Can llm already serve as a database interface? a big bench for large-scale database grounded text-to-sqls. *ArXiv*, abs/2305.03111.
- Aiwei Liu, Xuming Hu, Lijie Wen, and Philip S. Yu. 2023a. A comprehensive evaluation of chatgpt’s zero-shot text-to-sql capability. *ArXiv*, abs/2303.13547.
- Hu Liu, Yuliang Shi, Jianlin Zhang, Xinjun Wang, Hui Li, and Fanyu Kong. 2023b. Multi-hop relational graph attention network for text-to-sql parsing. *2023 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.
- Tom J Pollard, Alistair EW Johnson, Jesse D Raffa, Leo A Celi, Roger G Mark, and Omar Badawi. 2018. The eicu collaborative research database, a freely available multi-center database for critical care research. *Scientific data*, 5(1):1–13.
- Octavian Popescu, Irene Manotas, Ngoc Phuoc An Vo, Hangu Yeo, Elahe Khorashani, and Vadim Sheinin. 2022. Addressing limitations of encoder-decoder based approach to text-to-sql. In *International Conference on Computational Linguistics*.
- Jiexing Qi, Jingyao Tang, Ziwei He, Xiangpeng Wan, Chenghu Zhou, Xinbing Wang, Quanshi Zhang, and Zhouhan Lin. 2022. Rasat: Integrating relational structures into pretrained seq2seq model for text-to-sql. *ArXiv*, abs/2205.06983.
- Bowen Qin, Binyuan Hui, Lihan Wang, Min Yang, Jinyang Li, Binhua Li, Ruiying Geng, Rongyu Cao, Jian Sun, Luo Si, Fei Huang, and Yongbin Li. 2022. A survey on text-to-sql parsing: Concepts, methods, and future directions. *ArXiv*, abs/2208.13629.
- Preethi Raghavan, Jennifer J. Liang, Diwakar Mahajan, Rachita Chandra, and Peter Szolovits. 2021. emrbqa: A clinical knowledge-base question answering dataset. In *Workshop on Biomedical Natural Language Processing*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. *ArXiv*, abs/1806.03822.
- Jaydeep Sen, Chuan Lei, Abdul Quamar, Fatma Özcan, Vasilis Efthymiou, Ayushi Dalmia, Greg Stager, Ashish R. Mittal, Diptikalyan Saha, and Karthik Sankaranarayanan. 2020. Athena++. *Proceedings of the VLDB Endowment*, 13:2747 – 2759.

- Ruoxi Sun, Sercan Ö. Arik, Hootan Nakhost, Hanjun Dai, Rajarishi Sinha, Pengcheng Yin, and Tomas Pfister. 2023. [Sql-palm: Improved large language model adaptation for text-to-sql](#). *ArXiv*, abs/2306.00739.
- Richard Tarbell, Kim-Kwang Raymond Choo, Glenn Dietrich, and Anthony Rios. 2023. [Towards understanding the generalization of medical text-to-sql models and datasets](#). *AMIA ... Annual Symposium proceedings. AMIA Symposium*, 2023:669–678.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *ArXiv*, abs/2302.13971.
- Immanuel Trummer. 2022. [Codexdb: Synthesizing code for query processing from natural language instructions using gpt-3 codex](#). *Proc. VLDB Endow.*, 15:2921–2928.
- Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 2019. [Rat-sql: Relation-aware schema encoding and linking for text-to-sql parsers](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Kun Xu, Lingfei Wu, Zhiguo Wang, Mo Yu, Liwei Chen, and Vadim Sheinin. 2018. [Sql-to-text generation with graph-to-sequence model](#). *ArXiv*, abs/1809.05255.
- Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. [Tabert: Pretraining for joint understanding of textual and tabular data](#). *arXiv preprint arXiv:2005.08314*.
- Tao Yu, Rui Zhang, Kai-Chou Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Z Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir R. Radev. 2018. [Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task](#). *ArXiv*, abs/1809.08887.
- Xiang Yue, Bernal Jimenez Gutierrez, and Huan Sun. 2020. [Clinical reading comprehension: A thorough analysis of the emrqa dataset](#). *ArXiv*, abs/2005.00574.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. [Seq2sql: Generating structured queries from natural language using reinforcement learning](#). *ArXiv*, abs/1709.00103.

## Appendix A Examples of unanswerable questions

Category	Example	Frequency
Incorrect Patient Number	Will they have any urine test done for patient 24628?	252 (56%)
	Is patient 21074 subject to tests involving covid-19?	
	Do you know what type of blood patient 1903 has?	
Require External Knowledge	What is a checklist before lumb/lmbosac fus ant/ant?	83 (18.4%)
	What is the protocol used for the anticancer drugs?	
	So tell me what to do before you go for hemodialysis.	
Out of EHR Knowledge Base	What kind of blood patient 18866 has.	115 (25.6%)
	List the single rooms that are available now?	
	When are dr. oneill's rounds and procedures?	

Table 6: Examples of Unanswerable Questions with Respective Frequencies

## Appendix B Prompt

Task Description	
<p>I have a database related to healthcare that consists of 17 tables, each holding various pieces of data about hospital operations. I need to convert a natural language question into an SQL query to retrieve specific information from this database. Could you construct an SQL query that accurately reflects the question, considering the structure and details of my database? My database uses SQLite, and I'm looking for a query that's optimized for accuracy and efficiency. The detailed description of 17 tables is given below. Note that the patient number is an eight-digit number and current year is 2100 and all table values are in lower case.</p>	
Table Information – <Step 2> Initial SQL Generation	
Natural Language Form	SQL Form
<p>1.admissions: Documents each hospitalization event. • admissions [row_id, subject_id, hadm_id (hospital admission ID), admittance (admission time), disctime (discharge time), admission_type, admission_location, discharge_location, insurance, language, marital_status, age];</p> <p>2. d_icd_diagnoses: A reference for ICD-9 diagnosis codes • d_icd_diagnoses [row_id, icd_code, long_title];</p> <p>3. d_icd_procedures: A reference for ICD-9 procedure codes. • d_icd_procedures [row_id, icd_code, long_title];</p> <p>4. d_labitems: Acts as a dictionary for lab test ITEMIDs. • d_labitems [row_id, itemid, label]</p> <p>...</p>	<p>1.admissions: Documents each hospitalization event. <code>CREATE TABLE</code> admissions (   row_id <code>INTEGER</code>,   subject_id <code>INTEGER REFERENCES</code> patients(subject_id),   hadm_id <code>INTEGER</code>,   admittime <code>TIMESTAMP</code>,   disctime <code>TIMESTAMP</code>,   admission_type <code>TEXT</code>,   admission_location <code>TEXT</code>,   discharge_location <code>TEXT</code>,   insurance <code>TEXT</code>,   language <code>TEXT</code>,   marital_status <code>TEXT</code>,   age <code>INTEGER</code> );</p> <p>2. d_icd_diagnoses: A reference for ICD-9 diagnosis codes. <code>CREATE TABLE</code> d_icd_diagnoses (   row_id <code>INTEGER</code>,   icd_code <code>TEXT</code>,   long_title <code>TEXT</code> );</p> <p>3. d_icd_procedures: A reference for ICD-9 procedure codes. <code>CREATE TABLE</code> d_icd_procedures (   row_id <code>INTEGER</code>,   icd_code <code>TEXT</code>,   long_title <code>TEXT</code> );</p> <p>4. d_labitems: Acts as a dictionary for lab test ITEMIDs. <code>CREATE TABLE</code> d_labitems (   row_id <code>INTEGER</code>,   itemid <code>INTEGER</code>,   label <code>TEXT</code> );</p> <p>...</p>
Templatized Question	
<p>Count how many times patient 10004457 had 'coronary arteriography using two catheters' during their first hospital visit.</p>	

Table 7: The prompt used in the step 2 initial SQL generation.

## Appendix C Detailed Prompt Example

Task Description
<p>I have a database related to healthcare that consists of 17 tables, each holding various pieces of data about hospital operations. I need to convert a natural language question into an SQL query to retrieve specific information from this database. Could you construct an SQL query that accurately reflects the question, considering the structure and details of my database? My database uses SQLite, and I'm looking for a query that's optimized for accuracy and efficiency. The detailed description of 17 tables is given below. Note that the patient number is an eight-digit number and current year is 2100 and all table values are in lower case.</p>
Table Information – <Step 3> SQL Regeneration
<ol style="list-style-type: none"> <li>1. "procedures_icd": <ul style="list-style-type: none"> <li>• Records procedures using ICD codes.</li> <li>• procedures_icd (row_id, subject_id, hadm_id, icd_code, charttime)</li> <li>• Description of Columns: <ul style="list-style-type: none"> <li>row_id: Unique record identifier;</li> <li>subject_id: Unique identifier assigned to each individual patient;</li> <li>hadm_id (hospital admission ID): Unique identifier assigned to each separate hospital admission of a patient;</li> <li>icd_code: ICD code for the procedure performed;</li> <li>charttime: Timestamp of when the procedure was documented in the patient's chart",</li> </ul> </li> </ul> </li> <li>2. "d_icd_procedures": <ul style="list-style-type: none"> <li>• A reference for ICD-9 procedure codes.</li> <li>• d_icd_procedures (row_id, icd_code, long_title)</li> <li>• Description of Columns: <ul style="list-style-type: none"> <li>row_id: Unique record identifier;</li> <li>icd_code: Unique ICD-9 procedure code;</li> <li>long_title: Detailed description of the procedure.",</li> </ul> </li> </ul> </li> <li>3. "admissions": <ul style="list-style-type: none"> <li>• Documents each hospitalization event.</li> <li>• admissions (row_id, subject_id, hadm_id, admittance, dischtime, admission_type, admission_location, discharge_location, insurance, language, marital_status, age)</li> <li>• Description of Columns: <ul style="list-style-type: none"> <li>row_id: Unique record identifier;</li> <li>subject_id: Unique identifier assigned to each individual patient;</li> <li>hadm_id (hospital admission ID): Unique identifier assigned to each separate hospital admission of a patient;</li> <li>admittime: Admission time to the hospital;</li> <li>dischtime: Discharge time from the hospital;</li> <li>admission_type: Type of hospital admission;</li> <li>admission_location: Location from where the patient was admitted;</li> <li>discharge_location: Location to where the patient was discharged;</li> <li>insurance: Patient's insurance type;</li> <li>language: Patient's primary language;</li> <li>marital_status: Patient's marital status;</li> <li>age: Patient's age at admission.",</li> </ul> </li> </ul> </li> </ol>
Templatized Question
<p>Count how many times patient 10004457 had 'coronary arteriography using two catheters' during their first hospital visit.</p>

Table 8: The prompt used in Step 3 for SQL regeneration.

## Appendix D Word frequencies: 3-gram

N-gram	1-gram	2-gram	3-gram
<b>Train set</b>	('patient',): 3205 ( 'since',): 1572 ( 'last',): 1394 ( 'hospital',): 1340 ( 'first',): 1232 ( 'year',): 934 ( 'patients',): 861 ( '2100',): 818 ( 'visit',): 778 ( 'time',): 734	('hospital', 'visit'): 608 ( 'since', '2100'): 425 ( 'first', 'hospital'): 327 ( 'last', 'time'): 317 ( 'hospital', 'encounter'): 316 ( 'last', 'hospital'): 302 ( 'since', '1'): 302 ( '1', 'year'): 298 ( 'year', 'ago'): 298 ( 'first', 'time'): 280	('since', '1', 'year'): 298 ( '1', 'year', 'ago'): 298 ( 'first', 'hospital', 'visit'): 203 ( 'last', 'hospital', 'visit'): 183 ( 'within', '2', 'months'): 153 ( 'last', 'time', 'patient'): 118 ( 'first', 'time', 'patient'): 104 ( 'first', 'hospital', 'encounter'): 94 ( 'last', 'hospital', 'encounter'): 92 ( 'measured', 'last', 'hospital'): 92
<b>Dev set</b>	('patient',): 656 ( 'hospital',): 322 ( 'since',): 316 ( 'patients',): 279 ( 'last',): 255 ( 'first',): 253 ( 'year',): 202 ( 'visit',): 199 ( '2100',): 170 ( 'time',): 134	('hospital', 'visit'): 150 ( 'since', '2100'): 82 ( 'last', 'hospital'): 72 ( 'since', '1'): 70 ( '1', 'year'): 69 ( 'year', 'ago'): 69 ( 'first', 'time'): 67 ( 'hospital', 'encounter'): 65 ( 'first', 'hospital'): 59 ( 'lab', 'test'): 46	('since', '1', 'year'): 69 ( '1', 'year', 'ago'): 69 ( 'last', 'hospital', 'visit'): 46 ( 'first', 'hospital', 'visit'): 40 ( 'within', '2', 'months'): 30 ( 'last', 'hospital', 'encounter'): 20 ( 'first', 'time', 'patient'): 20 ( 'current', 'hospital', 'visit'): 19 ( 'arterial', 'blood', 'pressure'): 19 ( 'top', 'three', 'frequent'): 18
<b>Test set</b>	('patient',): 620 ( 'hospital',): 318 ( 'patients',): 306 ( 'since',): 304 ( 'last',): 265 ( 'first',): 234 ( 'year',): 217 ( '2100',): 184 ( 'visit',): 164 ( 'prescribed',): 140	('hospital', 'visit'): 128 ( 'since', '2100'): 100 ( 'hospital', 'encounter'): 81 ( 'first', 'hospital'): 72 ( 'last', 'hospital'): 72 ( 'since', '1'): 57 ( '1', 'year'): 57 ( 'year', 'ago'): 57 ( 'many', 'patients'): 54 ( 'number', 'patients'): 47	('since', '1', 'year'): 57 ( '1', 'year', 'ago'): 57 ( 'first', 'hospital', 'visit'): 42 ( 'last', 'hospital', 'visit'): 42 ( 'within', '2', 'months'): 41 ( 'last', 'hospital', 'encounter'): 26 ( 'measured', 'last', 'hospital'): 23 ( 'first', 'time', 'patient'): 20 ( 'first', 'hospital', 'encounter'): 20 ( 'last', 'time', 'patient'): 18
<b>Unanswerable Train set</b>	('patient',): 252 ( 'department',): 49 ( 'tell',): 42 ( 'procedure',): 41 ( 'blood',): 36 ( 'dr',): 36 ( 'received',): 34 ( 'rooms',): 29 ( 'test',): 28 ( 'last',): 27	('received', 'department'): 20 ( 'outpatient', 'schedule'): 18 ( 'rounds', 'procedures'): 17 ( 'another', 'department'): 16 ( 'rooms', 'available'): 15 ( 'diagnosis', 'patient'): 15 ( 'operating', 'rooms'): 14 ( 'blood', 'transfusion'): 14 ( 'name', 'diagnosis'): 14 ( 'genetic', 'test'): 14	('name', 'diagnosis', 'patient'): 12 ( 'last', 'time', 'patient'): 11 ( 'many', 'operating', 'rooms'): 11 ( 'appointment', 'another', 'department'): 11 ( 'genetic', 'test', 'patient'): 10 ( 'subject', 'covid-19', 'testing'): 9 ( 'type', 'blood', 'patient'): 9 ( 'ward', 'id', 'patient'): 9 ( "today's", 'outpatient', 'schedule'): 9 ( 'outpatient', 'schedule', 'dr'): 8

Table 9: 3-Gram frequency table with 10 examples sorted in non-increasing order

## Appendix E Examples of Question Templating

<b>Input: "When does patient 8016's influenza quarantine end?"</b>	
<b>Candidate templates</b>	"What was the time of <patient number>'s last influenza a/b by dfa microbiology test since 03/2100?"
	"Can you tell me when <patient number> had their first rapid respiratory viral screen & culture microbiology test in 08/this year?"
	"When did <patient number> depart hospital during this year for the last time?"
<b>Reformulated: "When is the end date of patient 8016's influenza quarantine?"</b>	
<b>Input: "Pull up the IDs of patients who were diagnosed with cataract extraction status."</b>	
<b>Candidate templates</b>	"Number of patients who were diagnosed with <unk>."
	"Number of patients who were diagnosed throughout this year with <unk>."
	"Tell me the number of patients diagnosed with <unk>."
<b>Reformulated: "Retrieve the ids of patients diagnosed with 'cataract extraction status.'"</b>	
<b>Input: "How many duloxetine prescription cases were there since 1 year ago?"</b>	
<b>Candidate templates</b>	"How much duloxetine has been prescribed to <patient number> in 05/2100 in total?"
	"How many drugs have been prescribed to <patient number> since 2 months ago?"
	"What is the number of drugs <patient number> was prescribed since 1 year ago?"
<b>Reformulated: "What is the number of duloxetine prescription cases since 1 year ago?"</b>	

Table 10: Examples of question templating using question masked templates.