

Uncertainty-Aware Text-to-Program for Question Answering on Structured Electronic Health Records

Daeyoung Kim
Seongsu Bae
Seungho Kim
Edward Choi

KAIST, Republic of Korea

DAEYOUNG.K@KAIST.AC.KR
SEONGSU@KAIST.AC.KR
SHOKIM@KAIST.AC.KR
EDWARDCHOI@KAIST.AC.KR

Abstract

Question Answering on Electronic Health Records (EHR-QA) has a significant impact on the healthcare domain, and it is being actively studied. Previous research on structured EHR-QA focuses on converting natural language queries into query language such as SQL or SPARQL (NLQ2Query), so the problem scope is limited to pre-defined data types by the specific query language. In order to expand the EHR-QA task beyond this limitation to handle multi-modal medical data and solve complex inference in the future, more primitive systemic language is needed. In this paper, we design the program-based model (NLQ2Program) for EHR-QA as the first step towards the future direction. We tackle MIMICSPARQL*, the graph-based EHR-QA dataset, via a program-based approach in a semi-supervised manner in order to overcome the absence of gold programs. Without the gold program, our proposed model shows comparable performance to the previous state-of-the-art model, which is an NLQ2Query model (0.9% gain). In addition, for a reliable EHR-QA model, we apply the uncertainty decomposition method to measure the ambiguity in the input question. We empirically confirmed data uncertainty is most indicative of the ambiguity in the input question.

Data and Code Availability Our source code and dataset are available on the official repository¹.

1. Introduction

Electronic health records (EHR) are composed of heterogeneous data (*e.g.*, medical history, diagnoses, radiology images, and test results) generated after patients receive some form of medical service. EHRs are

stored in databases with complex schemas such as in MIMIC-III (Johnson et al., 2016) or eICU (Pollard et al., 2018). It is hard for non-database experts to look up information or make decisions based on EHRs because they need to understand two things to obtain the information they want: the complex database structure and a query language such as SQL or SPARQL. For example, to answer the question “*what number of patients have been diagnosed with hyperglycemia?*”, one must generate a complex query such as “*select count (distinct patients.subject_id) from patients inner join admissions on patients.subject_id = admissions.subject_id where admissions.diagnosis = hyperglycemia*”. Therefore, a real-time QA agent that can understand the structure of EHRs and make complex inferences would significantly lower the burden of medical personnel during decision making, patients seeking information, and researchers conducting medical research.

Recent works on EHR-QA with structured data (*e.g.*, relational database or knowledge graph) have been focused on converting natural language questions (NLQ) into query languages such as SQL or SPARQL (Wang et al., 2020; Park et al., 2021; Bae et al., 2021) or into domain-specific forms (Raghavan et al., 2021). However, because all previous works mentioned above rely on specific query languages, the problem scope is limited to pre-defined data types (*e.g.*, string, int, timestamp) and operations. To expand the EHR-QA task beyond the scope of a query language in order to conduct more complex inference and use multiple modalities (*e.g.*, text, images, and signals), we require a program-based approach using atomic operations that are more primitive than those pre-defined by query languages. For instance, given a sufficiently powerful vision component, the program-based approach can answer questions such as “*Did*

1. <https://github.com/cyc1am3n/text2program-for-ehr>

this patient have effusion in the left lung before he was admitted to the ICU?”.

Within the healthcare domain, the reliability of deep neural network models is crucial because incorrect decisions bring consequences such as ethical issues on human life or monetary cost (Dusenberry et al., 2020). Likewise, EHR-QA models also need to be extremely reliable so that only correct answers are provided to the user, but the ambiguity in question due to lack of information or typos within the data make it challenging to train such reliable models. For example, if the word “procedure” is missing in the question in Figure 1, the question becomes ambiguous since “short title” could refer to either “procedure” or “diagnoses”, which increases uncertainty in data. Therefore, measuring uncertainty to detect ambiguous questions helps make a reliable model, as the model can take appropriate actions such as asking the user for clarification.

For the first time, our work uses the natural language question-to-program (NLQ2Program) approach for EHR-QA. Specifically, we tackle MIMIC-SPARQL* (Park et al., 2021), an EHR-QA dataset based on the open-source EHR data MIMIC-III. Since MIMICSPARQL* consists of pairs of a natural language question (NLQ) and a corresponding SPARQL query, all previous studies tackled this dataset by translating NLQ to either SQL or SPARQL queries with varying degrees of success. However, as stated above, we must venture beyond using a pre-defined query language such as SQL and SPARQL in order to handle multi-modal medical data and solve complex inference tasks in the future. Therefore we tackle MIMICSPARQL* via an NLQ2Program approach in a semi-supervised manner, in order to overcome the fact that there is no ground truth program given for each NLQ to train the model with. Our proposed model showed comparable performance to state-of-the-art NLQ2SQL or NLQ2SPARQL models that use the ground truth data. Also, we propose a method for measuring the ambiguity of input questions with insufficient information using the ensemble-based uncertainty decomposition for each program token generated by the EHR-QA model. We empirically demonstrate the effectiveness of using uncertainty decomposition to discern ambiguous questions, by evaluating MIMIC-SPARQL*’s test questions, where each question’s ambiguity was manually annotated.

The contributions we make in this paper can be summarized as follows:

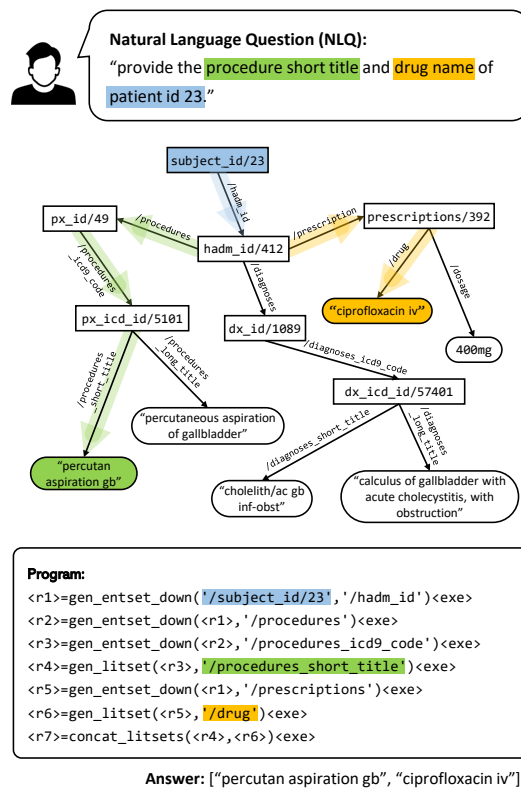


Figure 1: An illustrative example of our NLQ2Program approach for EHR question answering: a natural language question (NLQ), corresponding program traces over a knowledge graph, and its answer.

- It is the first attempt at designing an NLQ2Program model that uses programs composed of various atomic operations for an EHR-QA task. Without ground truth programs, we obtained results on MIMICSPARQL*, the most recent EHR-QA dataset, comparable to the NLQ2SQL and NLQ2SPARQL SOTA models that use ground truth queries (0.9% improvement).
- We generated a dataset to solve the problem without gold programs. We make it publicly available along with an interpreter that can execute programs so others can further research on EHR-QA using this NLQ2Program model in the future.

- We apply the ensemble-based uncertainty decomposition method to measure the ambiguity in the input question. To our best knowledge, this is the first attempt to detect ambiguous input questions in the QA research area. We show the effectiveness of measuring ambiguity using data uncertainty.

2. Related Works

2.1. QA on Electronic Health Record

Question answering on electronic health records (EHR-QA) can be divided into two broad categories: unstructured QA and structured QA. In the former case, most works focus on the machine reading comprehension task on free-formed text such as clinical case reports (Suster and Daelemans, 2018) and discharge summaries (Pampari et al., 2018). In the latter case, depending on database types of structured EHR, it can be further classified into two sub-categories: table-based QA and graph-based QA. In both subcategories, EHR-QA is treated as a translation task, converting a natural language question into a query language (*i.e.*, SQL/SPARQL) or a domain-specific logical form. Wang et al. (2020) first released MIMICSQL, a large-scale table-based EHR-QA dataset for the Question-to-SQL generation task in the healthcare domain, and also proposed a sequence-to-sequence (seq2seq) based model TREQS, which translates natural language questions to SQL queries (NLQ2SQL). Park et al. (2021) constructed a Question-to-SPARQL dataset and treated EHR-QA as a graph-based task by converting the original tables of MIMICSQL into a knowledge graph. Also, they empirically showed that NLQ2SPARQL outperforms NLQ2SQL for the same dataset and the same model architecture.

Recently, Raghavan et al. (2021) constructed a new large-scale question-logical form pair dataset (emrKBQA) for MIMIC-III, which reuses the same logical forms proposed in emrQA (Pampari et al., 2018), but it is not currently publicly available. Moreover, in order to execute the logical forms in emrKBQA, they must be mapped to corresponding SQL queries in advance. To overcome the limitation of a query language (*i.e.*, bound by pre-defined operations and only capable of handling fixed data types), we use the NLQ2Program approach for EHR-QA where programs are composed of atomic operations. Specifically, we develop our NLQ2Program approach while

viewing the EHR data as a knowledge graph rather than relational tables, similar to Park et al. (2021); Bae et al. (2021), where the graph-based approach outperformed the table-based approach.

2.2. Program Based Approach for KBQA

There are recent works translating natural language questions into multi-step executable programs over Knowledge Base Question Answering (KBQA; Structured QA) (Liang et al., 2017; Saha et al., 2019; Hua et al., 2020). These studies usually tackle datasets that do not have gold programs such as CQA (Saha et al., 2018) and WebQuestionSP (Yih et al., 2015). Specifically, Complex Question Answering (CQA) (Saha et al., 2018) is a large-scale QA dataset that contains complex questions involving multi-hop and aggregation questions (*e.g.*, counting, intersection, comparison), which are similar to our main target dataset (*i.e.*, MIMICSPARQL*). To handle the absence of the gold program, previous works proposed reinforcement learning (RL) based approaches.

RL-based approaches, however, face challenges caused by the large search space and sparse rewards. In particular, these challenges are intensified in MIMICSPARQL*, where KB artifacts (entities, relation, literal) are not explicitly revealed in the question. For example, for the NLQ in Figure 1, the QA model must generate a program using `‘/procedure_icd9_code’` and `‘prescriptions’` which were never mentioned in the NLQ. Moreover, MIMICSPARQL*’s search space is much larger than CQA since questions typically require a longer chain of operations to complete a program. Therefore instead of using RL, we train our model in a semi-supervised manner and compare our approach with NS-CQA (Hua et al., 2020), the state-of-the-art model for CQA.

2.3. Uncertainty in Language Generation

As the uncertainty in program-based EHR-QA has not been discussed before, we found the uncertainty in language generation to be the most relevant work to ours. Recent approaches focus on predictive uncertainty by measuring the probability (Ott et al., 2018) and the entropy (Xu et al., 2020; Xiao and Wang, 2021) of each token in the generated sequence by the model. Xiao and Wang (2021) apply the deep ensemble method (Lakshminarayanan et al., 2017) to decompose uncertainty into *data uncertainty*, the intrinsic uncertainty associated with

data, and *model uncertainty*², which reflects the uncertainty in model weights (Der Kiureghian and Ditlevsen, 2009; Kendall and Gal, 2017). These works analyze relationships between the uncertainty during decoding and the final output quality rather than analyzing the ambiguity in the input. Recently, Malinin and Gales (2020) utilize uncertainty for error detection and out-of-domain (OOD) input detection. For the OOD input detection, they focus on model uncertainty, which can capture discrepancies between the train and test datasets. However, our interest is to deal with insufficient information in the input question, which raises data uncertainty rather than model uncertainty.

3. Methodology

3.1. Preliminary: Dataset

In this work, we use a knowledge graph (KG) and questions from MIMICSPARQL* (Park et al., 2021) consisting of 10,000 question-SPARQL pairs that cover 9 tables³ of MIMIC-III (Johnson et al., 2016), an open-source ICU dataset. Note that MIMICSPARQL* was derived from MIMICSQL (Wang et al., 2020), a table-based EHR-QA dataset for MIMIC-III, consisting of 10,000 question-SQL pairs. In other words, MIMICSPARQL* has the same question as MIMICSQL, but the ground truth queries and database format are different. Also, note that each question in MIMICSPARQL* has two forms: template-based (machine-generated) form and natural (rephrased by medical domain experts) form.

3.2. Grammar

We newly define a grammar that can effectively explore the KG of MIMICSPARQL*. Since these are atomic operations that can be executed within KG, it is easy to handle multi-modality by expanding the grammar in the future. This uses a total of 7 data types which are either KG artifacts or basic data types. KG artifacts consist of *entSet* (set of entities), *rel* (relation), *lit* (literal), *litSet* (set of literals), *litSets* (tuple of *litSets*), and basic data types consist of *int* and *float*. Table 1 and Figure 2 describe the 14 operations we defined and their example. We consider

2. Data and model uncertainty are also called aleatoric and epistemic uncertainty.

3. Patients, Admissions, Diagnoses, Prescriptions, Procedures, Lab Results, Diagnosis Code Dictionary, Procedure Code Dictionary, Lab Code Dictionary

Saha et al. (2019)’s work as our starting point, but we modify the set of operations to make them more suitable for complex EHR-QA. For example, we add operations such as *maximum_litset*, *minimum_litset*, and *average_litset* since EHR-QA often requires calculations using numeric data found in the KG.

3.3. Problem Formulation

Our goal is to translate an EHR-related question into an executable program over KG. Assume an underlying programming language \mathcal{L} . Let us denote a given question by a sequence of tokens $Q = \{x_1, \dots, x_{|Q|}\}$ and the corresponding program $P \in \mathcal{L}$ can be represented as $P = \{y_1, \dots, y_{|P|}\}$. Our model aims to maximize the conditional probability $p(P|Q)$. Note that each question in MIMICSPARQL* has two forms, which are template-based (machine-generated) form and natural (rephrased by medical domain experts) form. We define the former as $Q_T = \{x_1, \dots, x_{|Q_T|}\}$ and the latter as $Q_N = \{x'_1, \dots, x'_{|Q_N|}\}$.

3.4. Synthetic Question-Program Generation

Since our method uses a custom set of operations (as described in Section 3.2), so our main obstacle to using NLQ2Program is the absence of gold programs (*i.e.*, sequences of operations) for questions in MIMICSPARQL*. We indirectly handle this problem by mass-generating pairs of MIMICSPARQL*-like questions and their corresponding programs. Based on our preliminary analysis of the template-based questions in the MIMICSPARQL* train dataset, we first create a list of templates (*e.g.*, what is the *RELATION* of *ENTITY*?) and question types (*e.g.*, retrieve question). Our analysis revealed that MIMICSPARQL* questions could be divided into total eight categories of question types. Then we generate synthetic question Q_{syn} and the corresponding program P_{syn} pairs in a form similar to Q_T . We sample program P_{syn} and corresponding question Q_{syn} by exploring KG while executing custom-defined operations for each of the eight question types. Since KG schema of MIMICSPARQL* is complex, even if the same question type is given, the pattern of generated synthetic program sequence varies greatly depending on the required KG artifact (*i.e.*, relation, entity). Using this method, we generate 30,000 (Q_{syn}, P_{syn}) pairs for each type, and a total of 168,574 pairs are used after excluding duplicate questions or ones that already exist in MIMICSPARQL*. Note that, it

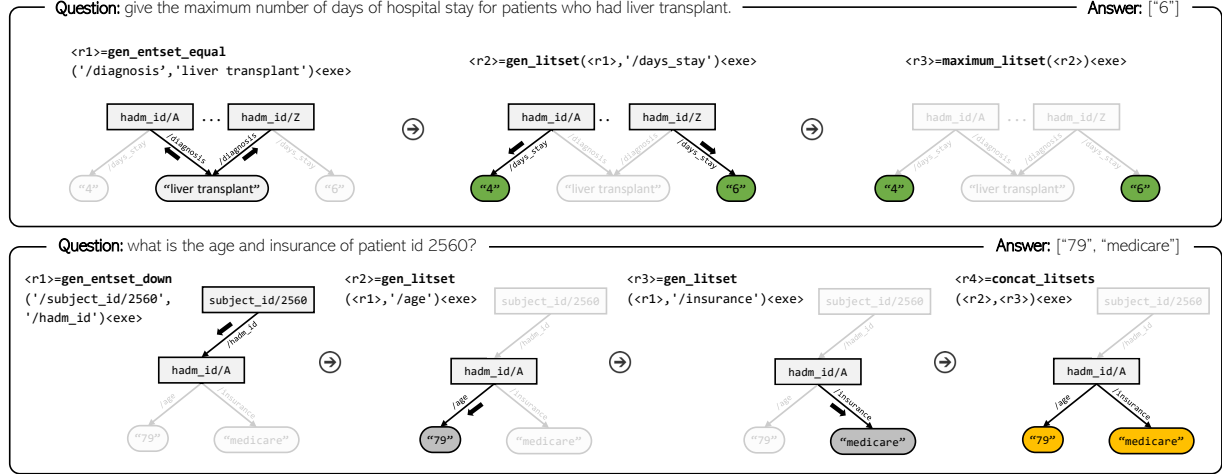


Figure 2: Illustrative examples of the natural language question (NLQ) and the corresponding program, composed of several predefined operations. To answer a single natural language question, we have to execute a series of atomic operations in sequence.

Table 1: Description of the custom-defined operations and return data types

Operation	Description	Return Data Type
<code>gen_entset_down(<i>entSet</i>, <i>rel</i>)</code>	the set of object entities associated with relation <i>rel</i> for each subject entity in the <i>entSet</i>	<i>entSet</i>
<code>gen_entset_up(<i>rel</i>, <i>entSet</i>)</code>	the set of subject entities associated with relation <i>rel</i> for each object entity in the <i>entSet</i>	<i>entSet</i>
<code>gen_litset(<i>entSet</i>, <i>rel</i>)</code>	the set of literal values associated with the relation <i>rel</i> for each subject entity of <i>entSet</i>	<i>litSet</i>
<code>gen_entset_equal(<i>rel</i>, <i>lit</i>)</code>	the set of subject entities which have literal value equal to <i>lit</i> for relation <i>rel</i>	<i>entSet</i>
<code>gen_entset_atleast(<i>rel</i>, <i>lit</i>)</code>	the set of subject entities which have literal value of at least <i>lit</i> for relation <i>rel</i>	<i>entSet</i>
<code>gen_entset_atmost(<i>rel</i>, <i>lit</i>)</code>	the set of subject entities which have literal value of at most <i>lit</i> for relation <i>rel</i>	<i>entSet</i>
<code>gen_entset_less(<i>rel</i>, <i>lit</i>)</code>	the set of subject entities which have smaller literal value than <i>lit</i> for relation <i>rel</i>	<i>entSet</i>
<code>gen_entset_more(<i>rel</i>, <i>lit</i>)</code>	the set of subject entities which have greater literal value than <i>lit</i> for relation <i>rel</i>	<i>entSet</i>
<code>count_entset(<i>entSet</i>)</code>	the number of entities in <i>entSet</i>	<i>int</i>
<code>intersect_entsets(<i>entSet</i>₁, <i>entSet</i>₂)</code>	the set of entities that exist in common in <i>entSet</i> ₁ and <i>entSet</i> ₂	<i>entSet</i>
<code>maximum_litset(<i>litSet</i>)</code>	the largest value in <i>litSet</i>	<i>float</i>
<code>minimum_litset(<i>litSet</i>)</code>	the smallest value in <i>litSet</i>	<i>float</i>
<code>average_litset(<i>litSet</i>)</code>	the average value of <i>litSet</i>	<i>float</i>
<code>concat_litsets(<i>litSet</i>₁, <i>litSet</i>₂)</code>	the combined list of <i>litSet</i> ₁ and <i>litSet</i> ₂	<i>litSets</i>

might be tempting to create gold programs by directly parsing the preexisting template questions Q_T , instead of creating synthetic questions and programs. This approach, however, has two major drawbacks: 1) it is complex; 2) Fragile to KG schema change. Further details about the synthetic data generation process are presented in the Appendix A. Note that there are a good number of questions unlikely to be asked in the real-world setting, because values are sampled by randomly exploring the KG (e.g., *how many patients whose language is engl and lab test value is 4.7k/ul?*).

3.5. Semi-supervised Learning

Obtaining NLQ & pseudo-gold program pairs

To acquire the pseudo-gold program \tilde{P} for the corresponding natural language question Q_N of MIMIC-SPARQL*, we introduce the following process:

1. Train a supplementary sequence-to-sequence model $f_{syn} : Q_{syn} \mapsto P_{syn}$ using synthetic pairs (i.e., question and program (Q_{syn}, P_{syn})).
2. Generate pseudo-gold programs \tilde{P} by feeding Q_T to the trained model f_{syn} in order to obtain corresponding (Q_T, \tilde{P}) pairs. Here, Q_T is used instead of Q_N since Q_{syn} and Q_T are based on the same templates.

3. Replace Q_T with Q_N to obtain pairs of natural-form questions and their corresponding pseudo-gold programs (Q_N, \tilde{P}).

Training Train a sequence-to-sequence model $f_N : Q_N \mapsto \tilde{P}$ with pairs of natural language question Q_N and pseudo-gold program \tilde{P} . Note that the synthetic pairs (Q_{syn}, P_{syn}) are not used for training f_N , and further experiments using synthetic pairs as pre-training data are shown in Appendix C.

3.6. Measuring Ambiguity of Question

As mentioned above, we can use uncertainty in the output program to detect ambiguous questions that lack essential information (*e.g.*, a user does not define the patient ID when asking for a patient’s age) or include unseen values (*e.g.*, typos). Typically, uncertainty can be divided into data uncertainty and model uncertainty, where the former can be viewed as uncertainty measuring the noise inherent in given training data, and the latter as uncertainty regarding noise in the deep neural network parameters (Chang et al., 2020; Dusenberry et al., 2020). Assuming that we can view ambiguous questions as inherent noise in the data (which the model cannot overcome by collecting more data, unlike model uncertainty), we aim to detect ambiguous input by measuring data uncertainty. Following Xiao and Wang (2021); Malinin and Gales (2020), we adopt the ensemble-based uncertainty estimation method.

Given the question $Q = \{x_1, \dots, x_{|Q|}\}$ and the corresponding program $P \in \mathcal{L}$, we denote the context of the i -th program token y_i as $c_i = \{x_1, \dots, x_{|Q|}, y_1, \dots, y_{i-1}\}$, the prediction of each model in the ensemble of M models as $\{p_m(y_i|c_i)\}_{m=1}^M$, and the aggregated prediction as $p(y_i|c_i) = \frac{1}{M} \sum_{m=1}^M p_m(y_i|c_i)$. Given context c_i , the entropy of $p_m(y_i|c_i)$ and $p(y_i|c_i)$ can be calculated as follows:

$$H_m(y_i|c_i) = - \sum_{v \in \mathcal{V}} p_m(y_i = v|c_i) \log p_m(y_i = v|c_i)$$

$$H(y_i|c_i) = - \sum_{v \in \mathcal{V}} p(y_i = v|c_i) \log p(y_i = v|c_i)$$

where \mathcal{V} is the whole vocabulary. $H(y_i|c_i)$ represents the total uncertainty which is sum of data and model uncertainty. Then we can decompose $H(y_i|c_i)$ into

data and model uncertainty as follows:

$$u_{\text{data}}(y_i|c_i) = \frac{1}{M} \sum_{m=1}^M H_m(y_i|c_i)$$

$$u_{\text{model}}(y_i|c_i) = H(y_i|c_i) - u_{\text{data}}(y_i|c_i)$$

We assume the ambiguity of NLQ will raise the data uncertainty of a specific program token, not the entire program itself. Note that $u_{\text{data}}(y_i|c_i)$ is calculated for every program token y_i . We utilize the maximum value of the data uncertainty $u_{\text{data}}(y_i|c_i)$ for every program token y_i , instead of aggregating the $u_{\text{data}}(y_i|c_i)$ in a program-level manner (Malinin and Gales, 2020). Specifically, we determine if the input question is ambiguous using detector g as follows:

$$g(\mathbf{U}; \tau) = \begin{cases} 0 & \text{if } \max(\mathbf{U}) \leq \tau \\ 1 & \text{if } \max(\mathbf{U}) > \tau \end{cases}$$

where $\mathbf{U} = \{u_{\text{data}}(y_1|c_1), \dots, u_{\text{data}}(y_{|P|}|c_{|P|})\}$ and specific threshold τ . We empirically show that this method can effectively detect ambiguous input questions.

4. Experiments

4.1. Experiment Settings

4.1.1. MODEL CONFIGURATIONS

Both the pseudo-gold program generating model f_{syn} and the NLQ2Program model f_N can be initialized with any sequence-to-sequence structure. We choose T5-base (Raffel et al., 2020), which is known to perform well in the natural language generation (NLG) field, for both models. For comparison, we also experiment with UniQA (Bae et al., 2021), the state-of-the-art model in the MIMICSPARQL* dataset. For the decoding strategy to generate program traces, we use beam search (Wiseman and Rush, 2016). Of the 8,000 samples from the MIMICSPARQL* training dataset, a total of 7,472 (Q_N, \tilde{P}) pairs that return the same execution result as the ground truth SPARQL query are used. For an accurate evaluation, we use 949 of 1,000 samples from the MIMICSPARQL* test dataset after excluding samples whose ground truth SPARQL execution returns NULL, or whose questions and SPARQL queries do not match (*e.g.*, the question adds the condition “less than 60 years of age” while the ground truth query looks for “DEMOGRAPHIC.AGE < 62”).

Table 2: Test results on MIMICSPARQL* with two different approaches: NLQ2SPARQL and NLQ2Program. We report the mean and standard deviation of execution accuracy (Acc_{EX}) over 5 random seeds.

Recovery Technique	NLQ2SPARQL (w/ ground truth query)				NLQ2Program (w/o gold program)			
	Seq2Seq	TREQS	UniQA	T5	NS-CQA (1%)	NS-CQA (100%)	Ours (UniQA)	Ours (T5)
✗	0.327 (0.043)	0.699 (0.013)	0.899 (0.010)	0.905 (0.006)	0.203 (0.043)	0.734 (0.087)	0.896 (0.009)	0.899 (0.005)
✓	0.338 (0.045)	0.712 (0.011)	0.939 (0.005)	0.937 (0.006)	-	-	0.938 (0.006)	0.948 (0.006)

4.1.2. BASELINES

In the experiment, we compare our model against five baseline models as follows: Seq2Seq (Luong et al., 2015), TREQS (Wang et al., 2020), UniQA (Bae et al., 2021), T5 (Raffel et al., 2020) and NS-CQA (Hua et al., 2020). The first four are NLQ2SPARQL models using ground truth SPARQL queries, and the last is an NLQ2Program model which does not use gold programs. Note that the T5 and UniQA models used as baselines adopt the NLQ2SPARQL approach, not NLQ2Program. Among the previous program-based approaches mentioned in Section 2.2, we choose NS-CQA as the baseline since it is the state-of-the-art model in the CQA (Saha et al., 2018) dataset. All models are trained with five random seeds, and we report the mean and standard deviation of performance. The details of implementation are provided in Appendix B.

Seq2Seq with Attention (NLQ2SPARQL)

Seq2Seq with attention (Luong et al., 2015) consists of a bidirectional LSTM encoder and an LSTM decoder. Following the original paper, we apply the attention mechanism in this model. Note that this model cannot handle the out-of-vocabulary (OOV) tokens. We denote the model as Seq2Seq.

TREQS (NLQ2SPARQL)

TREQS (Wang et al., 2020) is an LSTM-based encoder-to-decoder model using an attentive-copying mechanism and a recovery technique to handle the OOV problem.

UniQA (NLQ2SPARQL)

UniQA (Bae et al., 2021) is the state-of-the-art NLQ2Query model on MIMICSPARQL*. UniQA consists of a unified encoder-as-decoder architecture, which uses masked language modeling in the NLQ part and sequence-to-sequence modeling in the query part at the same time. Following the original paper,

we initialize UniQA with pre-trained BERT (12-layer, 768-hidden, 12-head) (Devlin et al., 2018).

T5 (NLQ2SPARQL)

T5 (Raffel et al., 2020) is a transformer-based encoder-to-decoder model which is pre-trained on a large corpus to convert every language problem into a text-to-text format. We use T5-base model (12-layer, 768-hidden, 12-head) as mentioned above.

NS-CQA (NLQ2Program)

NS-CQA (Hua et al., 2020) is an LSTM-based encoder-to-decoder RL framework that obtains state-of-the-art performance on the CQA dataset. It uses the copy mechanism and a masking method to reduce search space. A memory buffer, which stores promising trials for calculating a bonus reward, is used to alleviate the sparse reward problem. The model needs to be pre-trained by teacher forcing with pseudo-gold programs in order to mitigate the cold start problem. We pre-train the model with two different data settings to study the effectiveness of the RL approach with restricted semi-supervision: (1) pre-train using all (Q_N, \hat{P}) pairs, and (2) pre-train using only 1% of all (Q_N, \hat{P}) pairs (the same setting as (Hua et al., 2020)). We then fine-tune the model by employing RL using all of the Q_N and execution results of gold SPARQL query pairs.

4.1.3. EVALUATION METRIC

For comparing various models, three metrics are used in previous studies (Wang et al., 2020; Park et al., 2021; Bae et al., 2021), which are *Logical Form Accuracy* (Acc_{LF}), *Execution Accuracy* (Acc_{EX}), and *Structural Accuracy* (Acc_{ST}), to evaluate the generated queries (*i.e.*, SQL, SPARQL). However, Acc_{LF} and Acc_{ST} require gold programs since they compare the generated queries with the ground truth queries token by token. Therefore, we only use *Execution Accuracy* (Acc_{EX}), which measures the correctness

of the answer retrieved by executing the generated program with the KG⁴.

4.1.4. RECOVERING CONDITION VALUES

Following Wang et al. (2020), we apply the recovery technique to handle inaccurately generated condition values that often contain complex medical terminology. This technique replaces the condition values in the generated program with the most similar values that exist in the database. For instance, the user may ask “*how many patients had physical restrain status?*”, then one operation in the generated program could be “*gen_entset_equal('/diagnoses_long_title', 'physical_restrain_status')*”. However, in the database, the value ‘*physical_restraints_status*’ exists, but ‘*physical_restrain_status*’ does not. In that case, the recovery technique replaces the incorrect condition value of the program to the correct one, thus making it executable. In order to calculate the similarity between predicted values and existing ones, this technique uses ROUGE-L (Lin, 2004) score.

4.2. Experiment Result

As shown in Table 2, despite the absence of gold programs, our NLQ2Program model is comparable with state-of-the-art NLQ2SPARQL models that require ground truth query data. Also, we show that using NS-CQA’s semi-supervised RL-based approach on MIMICSPARQL* is not effective when using only 1% of the question and pseudo-gold program pairs for pre-training.

Note that the recovery technique is unnecessary for NS-CQA because, due to the decoding nature of NS-CQA, all KB artifacts (entity, relation, value) are copy-pasted from the NLQ to their appropriate locations after the program is generated.

Additionally, we conducted experiments regarding the effect of using the synthetic data introduced in Section 3.4, as pre-training data. The detailed information is shown in Appendix C.

4.3. Ambiguous Question Detection

In order to validate our method of measuring ambiguity using data uncertainty, we hand-annotated all MIMICSPARQL* test samples with the following labels. According to the degree of ambiguity, we categorized the ambiguous questions into two types: (1)

4. This is why we excluded test samples whose answers are NULL, to minimize lucky guesses.

mildly ambiguous (Mild) and (2) *highly ambiguous* (High). There are a total of 152 questions labeled as mildly ambiguous and 72 questions as highly ambiguous. The rules of ambiguous question labeling are defined as follows:

- **Mildly Ambiguous (Mild):** We treat questions whose relations are not explicitly revealed in the NLQ, so it is hard to infer even with the condition value as mildly ambiguous questions. For instance, it is challenging to know whether the relation to “*pneumococcal pneumonia*” is a */diagnoses_long_title* or a */diagnoses_short_title* for the question “*provide the number of patients less than 83 years of age who were diagnosed with pneumococcal pneumonia.*”. However, if there is an ideal linker connected to the DB, the ideal linker knows that “*pneumococcal pneumonia*” refers to a *short_title*. For this reason, we label these questions as mildly ambiguous.

In addition, typos in the questions are unseen values for the EHR-QA model that also cause ambiguity, so questions with typos are also labeled as mildly ambiguous (e.g., *what is the number of (dead →) cead patients who had brain mass; intracranial hemorrhage?*).

- **Highly Ambiguous (High):** We consider a question as highly ambiguous when its NLQ is too vague that multiple correct programs can be generated. For the question “*specify primary disease and icd9 code of patient id 18480*”, “*icd9_code*” could refer to either “*procedure*” or “*diagnosis*”.

In addition, if the condition value in an NLQ corresponds to more than one relation, that NLQ is labeled as highly ambiguous. Note that even an ideal linker cannot find the exact relation of highly ambiguous questions. For example, the question “*give the number of newborns who were born before the year 2168.*” is highly ambiguous since the condition value “*newborn*” is related to both “*/admission_type*” and “*/diagnosis*”.

We assess the detection capabilities of ambiguous questions based on data uncertainty u_{data} , model uncertainty u_{model} , total uncertainty H , and entropy H_m . These types of uncertainty are calculated from all tokens in a single program sequence and we use the maximum value as the degree of uncertainty. Note that u_{data} , u_{model} , and H are calculated via ensemble model consisting of 5 models, but H_m is calcu-

Table 3: Uncertainty results on MIMICSPARQL* with two ambiguity degrees: Mild & High and High. We report AUPRC and AUROC for four types of uncertainty: u_{data} , u_{model} , H , and H_m . Note that u_{data} , u_{model} , and H are calculated by the ensemble model consisting of 5 models, but H_m is calculated by a single model. In case of H_m , we report the mean and standard deviation over 5 random seeds.

Ambiguity		u_{data}	u_{model}	H	H_m
Mild & High	AUPRC	0.417	0.340	0.387	0.375 (0.02)
	AUROC	0.705	0.630	0.631	0.650 (0.02)
High	AUPRC	0.173	0.117	0.153	0.153 (0.01)
	AUROC	0.696	0.614	0.650	0.651 (0.01)

lated by a single model. Performance is assessed via the area under a Precision-Recall (PR) curve and a receiver operating characteristic (ROC) curve. The results in Table 3 show that using only data uncertainty detects ambiguous questions better than other metrics. However, ambiguous questions cannot be perfectly detected because ambiguous questions are also included in the training data, and we only labeled ambiguous questions in the test samples. Additionally, we compare the token-level and program-level methods of measuring uncertainty. The results show the advantage of the token-level method. Details of this experiment are available in Appendix D.

To implement a more practical QA system, an interface is required to interact with the user, asking clarifying questions or allowing the user to modify the generated program. Solving these issues, however, is beyond the scope of this work. Instead, we build the system that gives users all five best beam hypotheses (*i.e.*, programs) so users can select the appropriate candidate program when the ambiguity of the input question exceeds the specified threshold. The results in Figure 3 show that the execution accuracy increases up to 0.986 with the number of recommendations.

4.4. Qualitative Results

In this section, we provide qualitative results to analyze generated programs along with token-level data uncertainty.




Figure 3: Test results of MIMICSPARQL* for five single models and an ensemble model. For the ensemble model, all tokens are generated by aggregating the prediction of all single models. The upper part of the shaded area presents the maximum execution accuracy of a single model, and the lower part shows minimum accuracy.

4.4.1. GENERATED PROGRAMS FOR AMBIGUOUS QUESTIONS

As we expect, if the input question is ambiguous, high data uncertainty is measured. The question in the first example in Figure 4 corresponds to a highly ambiguous question since whether the relation is either ‘*long_title*’ or ‘*short_title*’ is not specified. It leads to the high uncertainty of the token ‘*s*’ whose data uncertainty is almost 100 times larger than the average data uncertainty of other tokens. Note that if the model generates ‘*long_title*’ rather than ‘*short_title*’, the execution result is incorrect but the program is semantically aligned with the question. Likewise, in the second example, the question corresponds to a mildly ambiguous question since the ‘*icd9 code 9229*’ exists only in the ‘*procedures.*’ However, it is hard to recognize this fact for the model, which does not have an ideal linker, so the data uncertainty is increased at the position of the token ‘*pro.*’ These examples also demonstrate the max value of token-level uncertainty is an effective representative of the ambiguity.

4.4.2. FAILURE CASES

There are some failure cases with high data uncertainty, but the NLQ is not ambiguous. In the first example of Figure 5, the token with the highest data uncertainty is “*gen_entset_down,*”, an incorrect operation that must be changed to “*gen_litset.*” Simi-



Figure 4: Qualitative results for ambiguous question. We visualize token-level data uncertainties of the generated program using the heatmap.

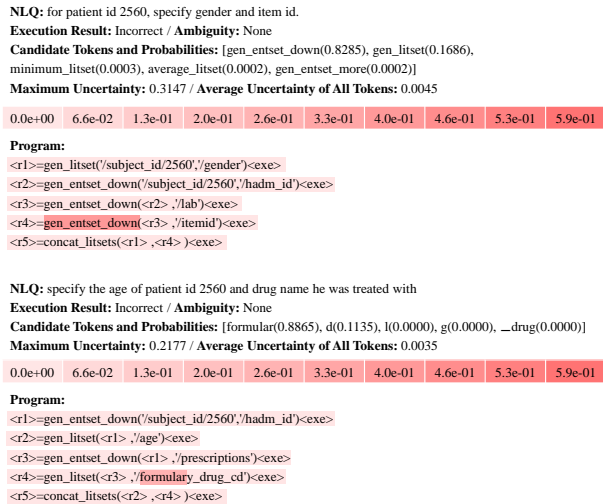


Figure 5: Qualitative results for failure cases. We visualize token-level data uncertainties of the generated program using the heatmap.

larly, in the second example, the token with the highest data uncertainty is “formular,” a subset of relation “/formulary_drug_cd” (drug code), which should

be changed to “/drug” (drug name). It can be seen that the data uncertainty does not always represent the ambiguity of the question. However, when the model generates uncertain tokens, interaction with the user can still help improve the performance and reliability of the EHR-QA model. Uncertainty in NLQ2Program is just the beginning, so more research is needed.

5. Conclusion

In this work, we designed an NLQ2Program methodology using atomic operations for EHR-QA task on MIMICSPARQL*. We tackled the absence of gold programs via NLQ2Program approach in a semi-supervised manner. Our proposed model showed comparable performance with the previous NLQ2SPARQL state-of-the-art model. Moreover, we applied the ensemble-based uncertainty decomposition method to detect the ambiguous input question. We showed the effectiveness of measuring ambiguity using data uncertainty. Our further direction is to extend our methodology to handle multi-modal data on EHR and solve more complex questions.

Institutional Review Board (IRB)

This research does not require IRB approval.

Acknowledgments

This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant (No.2019-0-00075, Artificial Intelligence Graduate School Program(KAIST)) and National Research Foundation of Korea (NRF) grant (NRF-2020H1D3A2A03100945), funded by the Korea government (MSIT).

References

Seongsu Bae, Daeyoung Kim, Jiho Kim, and Edward Choi. Question answering for complex electronic health records database using unified encoder-decoder architecture. In *Machine Learning for Health*, pages 13–25. PMLR, 2021.

Jie Chang, Zhonghao Lan, Changmao Cheng, and Yichen Wei. Data uncertainty learning in face

- recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5710–5719, 2020.
- Armen Der Kiureghian and Ove Ditlevsen. Aleatory or epistemic? does it matter? *Structural safety*, 31(2):105–112, 2009.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Michael W Dusenberry, Dustin Tran, Edward Choi, Jonas Kemp, Jeremy Nixon, Ghassen Jerfel, Katherine Heller, and Andrew M Dai. Analyzing the role of model uncertainty for electronic health records. In *Proceedings of the ACM Conference on Health, Inference, and Learning*, pages 204–213, 2020.
- Yuncheng Hua, Yuan-Fang Li, Guilin Qi, Wei Wu, Jingyao Zhang, and Daiqing Qi. Less is more: Data-efficient complex question answering over knowledge bases. *Journal of Web Semantics*, 65: 100612, 2020.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-Wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.
- Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- Chen Liang, Jonathan Berant, Quoc Le, Kenneth Forbus, and Ni Lao. Neural symbolic machines: Learning semantic parsers on freebase with weak supervision. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 23–33, 2017.
- Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, 2015.
- Andrey Malinin and Mark Gales. Uncertainty estimation in autoregressive structured prediction. In *International Conference on Learning Representations*, 2020.
- Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. Analyzing uncertainty in neural machine translation. In *International Conference on Machine Learning*, pages 3956–3965. PMLR, 2018.
- Anusri Pampari, Preethi Raghavan, Jennifer Liang, and Jian Peng. emrqa: A large corpus for question answering on electronic medical records. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2357–2368, 2018.
- Junwoo Park, Youngwoo Cho, Haneol Lee, Jaegul Choo, and Edward Choi. Knowledge graph-based question answering with electronic health records. In *Machine Learning for Healthcare Conference*, pages 36–53. PMLR, 2021.
- Tom J Pollard, Alistair EW Johnson, Jesse D Raffa, Leo A Celi, Roger G Mark, and Omar Badawi. The eicu collaborative research database, a freely available multi-center database for critical care research. *Scientific data*, 5(1):1–13, 2018.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67, 2020.
- Preethi Raghavan, Jennifer J Liang, Diwakar Mahajan, Rachita Chandra, and Peter Szolovits. emrkbqa: A clinical knowledge-base question answering dataset. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 64–73, 2021.

Amrita Saha, Vardaan Pahuja, Mitesh M Khapra, Karthik Sankaranarayanan, and Sarath Chandar. Complex sequential question answering: Towards learning to converse over linked question answer pairs with a knowledge graph. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Amrita Saha, Ghulam Ahmed Ansari, Abhishek Laddha, Karthik Sankaranarayanan, and Soumen Chakrabarti. Complex program induction for querying knowledge bases in the absence of gold programs. *Transactions of the Association for Computational Linguistics*, 7:185–200, 2019.

Simon Suster and Walter Daelemans. Clicr: a dataset of clinical case reports for machine reading comprehension. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1551–1563, 2018.

Ping Wang, Tian Shi, and Chandan K Reddy. Text-to-sql generation for question answering on electronic medical records. In *Proceedings of The Web Conference 2020*, pages 350–361, 2020.

Sam Wiseman and Alexander M Rush. Sequence-to-sequence learning as beam-search optimization. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1296–1306, 2016.

Yijun Xiao and William Yang Wang. On hallucination and predictive uncertainty in conditional language generation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2734–2744, 2021.

Jiacheng Xu, Shrey Desai, and Greg Durrett. Understanding neural abstractive summarization models via uncertainty. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6275–6281, 2020.

Scott Wen-tau Yih, Ming-Wei Chang, Xiaodong He, and Jianfeng Gao. Semantic parsing via staged query graph generation: Question answering with knowledge base. In *Proceedings of the Joint Conference of the 53rd Annual Meeting of the ACL and the 7th International Joint Conference on Natural Language Processing of the AFNLP*, 2015.

Appendix A. Details for Synthetic Data Generation

Based on our preliminary analysis of the template-based questions in the MIMICSPARQL* machine-generated train dataset, we found there are eight basic templates as follows:

- what is *RELATION* of *ENTITY*?
- what is *RELATION1* and *RELATION2* of *ENTITY*?
- what is *RELATION1* of *RELATION2 VALUE*?
- what is *RELATION1* and *RELATION2* of *RELATION3 VALUE*?
- what is the number of *ENTITY* whose *RELATION CONDITION LITERAL*?
- what is the number of *ENTITY* whose *RELATION1 CONDITION1 LITERAL1* and *RELATION2 CONDITION2 LITERAL2*?
- what is *AGGR RELATION1* of *ENTITY* whose *RELATION2 CONDITION LITERAL*?
- what is *AGGR RELATION1* of *ENTITY* whose *RELATION2 CONDITION1 LITERAL1* and *RELATION3 CONDITION2 LITERAL2*?

where *CONDITION* corresponds to $=, >, <, \leq, \geq$ and *AGGR* represents min, max, and average. For each template, we composed the operation set to be executed. When we generate the synthetic question and corresponding synthetic program, the operation and corresponding argument to be selected for each step are arbitrarily determined. Note that this simple technique is available for other KGs.

Appendix B. Implementation Details

We implement our model and baseline models with PyTorch Lightning⁵ and HuggingFace’s transformers⁶. In the case of TREQS and NS-CQA, we utilized the official code^{7,8} written by the origin authors. Also in the case of UniQA, we manually implement model followed by descriptions in Bae et al. (2021) since the official codes are not publicly available.

5. <https://www.pytorchlightning.ai>

6. <https://huggingface.co/transformers/>

7. <https://github.com/wangpinggl/TREQS>

8. <https://github.com/DevinJake/NS-CQA/>

B.1. Hyperparameters

In order to make an accurate comparison with the baseline models, the Seq2Seq model and TREQS model were imported from Park et al. (2021), and hyperparameters were also imported with the same value. We trained our models on the following GPU environment: NVIDIA GeForce RTX-3090. Also, torch version is 1.7.0, and CUDA version is 11.1. Other hyperparameters are presented in Table 4.

Appendix C. Performance Variance by Pre-trained Model

In Section 3.4, we introduced our synthetic data generation method via preliminary analysis of template-based questions in MIMICSPARQL* training set, and generated 168,574 synthetic question-program pairs. With this large volume of pairs, we can utilize them as source data for further pre-training a model to improve its final performance. Therefore, we investigate the utility of synthetic pairs for pre-training with three different types of models. As shown in Table 5, we can observe that three models, first initialized with BERT, improve after further pre-training with synthetic pairs. However, our synthetic data cannot give performance gain because T5 is originally pre-trained with massive data so that our synthetic data cannot give model performance gain with the same quality as the original corpus of T5-base.

Appendix D. Token-level vs. Program-level Uncertainty Measuring

We conduct the additional experiment to compare the token-level uncertainty measuring method with the program-level uncertainty measuring method. Following Malinin and Gales (2020), we utilize the import weighting method using all beam hypotheses. We use the normalizing factor for import weighting as 1 followed by the original paper. We denote the data uncertainty, the model uncertainty, and the total uncertainty in the program-level as U_{data} , U_{model} , and U_{total} respectively. Specially, U_{data} is obtained by subtract U_{model} from U_{total} . The result shows that the token-level uncertainty measuring method is more effective than program-level when detecting ambiguous questions. As we mentioned in Section 3.6, we speculate uncertainty only increases at the missing information part as shown in Section 4.4. In addition,

Table 4: Hyperparameters for training several models.

Hyperparameters	Seq2Seq	TREQS	UniQA	NS-CQA	Ours
Hidden dimension	256	256(enc) + 256(dec)	768	128	768
Learning rate	5×10^{-4}	5×10^{-4}	3×10^{-5}	1×10^{-3} (PT), 1×10^{-4} (RL)	1×10^{-4}
LR Scheduler	StepLR(step size = 2, step decay = 0.8)	StepLR(step size = 2, step decay = 0.8)	Linear decay	Linear decay	Linear decay
Batch size	16	64	32	32 (PT), 8 (RL)	32
Epochs	20	20	100 (w/ early stop)	100 (PT), 30 (RL) (w/ early stop)	100 (w/ early stop)
Seed	1, 12, 123, 1234, 42	1, 12, 123, 1234, 42	1, 12, 123, 1234, 42	1, 12, 123, 1234, 42	1, 12, 123, 1234, 42
Beam size	-	5	5	-	5

Table 5: Test results of four pre-trained models on MIMICSPARQL* depending on whether each model utilizes synthetic data (P_{syn} , Q_{syn}) or not. We report mean and standard deviation of execution accuracy Acc_{EX} over five random seeds.

Recovery Technique	Synthetic Data (P_{syn} , Q_{syn})	E-as-D (109M)	UniQA (109M)	BERT2BERT (130M)	T5-base (220M)
✗	✗	0.877 (0.013)	0.860 (0.015)	0.854 (0.006)	0.899 (0.005)
✗	✓	0.882 (0.003)	0.896 (0.009)	0.893 (0.003)	0.896 (0.006)
✓	✗	0.939 (0.009)	0.920 (0.010)	0.913 (0.009)	0.948 (0.006)
✓	✓	0.944 (0.008)	0.938 (0.006)	0.940 (0.004)	0.944 (0.005)

Table 6: Token/Program-level uncertainty results for two different ambiguity degrees: Mild & High and High. We report AUPRC and AUROC as evaluation metrics.

Ambiguity		Token-level			Program-level		
		u_{data}	u_{model}	H	U_{data}	U_{model}	U_{total}
Mild & High	AUPRC	0.417	0.340	0.387	0.325	0.141	0.144
	AUROC	0.705	0.630	0.631	0.626	0.179	0.195
High	AUPRC	0.173	0.117	0.153	0.129	0.051	0.052
	AUROC	0.696	0.614	0.650	0.633	0.312	0.328

the result demonstrates the fact the data uncertainty is most indicative of the ambiguity in the input question consistently.