

Topic Transferable Table Question Answering

Saneem A. Chemmengath^{1*} Vishwajeet Kumar^{1*} Samarth Bharadwaj¹
Jaydeep Sen¹ Mustafa Canim¹, Soumen Chakrabarti²
Alfio Gliozzo¹ Karthik Sankaranarayanan¹,

¹IBM Research ²Indian Institute of Technology, Bombay

saneem.cg@in.ibm.com, vishk024@in.ibm.com, samarth.b@in.ibm.com,
jaydesen@in.ibm.com, mustafa@us.ibm.com, soumen.chakrabarti@gmail.com
gliozzo@us.ibm.com, kartsank@in.ibm.com

Abstract

Weakly-supervised table question-answering (TableQA) models have achieved state-of-art performance by using pre-trained BERT transformer to jointly encoding a question and a table to produce structured query for the question. However, in practical settings TableQA systems are deployed over table corpora having topic and word distributions quite distinct from BERT’s pretraining corpus. In this work we simulate the practical topic shift scenario by designing novel challenge benchmarks WikiSQL-TS and WikiTQ-TS¹, consisting of train-dev-test splits in five distinct topic groups, based on the popular WikiSQL and WikiTableQuestions datasets. We empirically show that, despite pre-training on large open-domain text, performance of models degrades significantly when they are evaluated on unseen topics. In response, we propose T3QA (Topic Transferable Table Question Answering) a pragmatic adaptation framework for TableQA comprising of: (1) topic-specific vocabulary injection into BERT, (2) a novel text-to-text transformer generator (such as T5, GPT2) based natural language question generation pipeline focused on generating topic specific training data, and (3) a logical form re-ranker. We show that T3QA provides a reasonably good baseline for our topic shift benchmarks. We believe our topic split benchmarks will lead to robust TableQA solutions that are better suited for practical deployment.

1 Introduction

Documents, particularly in enterprise settings, often contain valuable tabular information (e.g., financial, sales/marketing, HR). Natural language question answering systems over a table (or TableQA) have an additional complexity of understanding the tabular structure including row/column head-

*Equal contribution by first two authors.

¹The source code and new dataset splits are available at <https://github.com/IBM/T3QA>

Party	Candidate	Votes
Conservatives	Andrew Turner	32,717
Liberal Democrats	Anthony Rowlands	19,739
Labour	Mark Chiverton	11,484
UK Independence	Michael Tarrant	2,352
Independent	Edward Corby	551

Question: Who ran in the election for labour party?

Answer : Mark Chiverton

Figure 1: Topic-sensitive representations are important to infer that, in the context of the topic *politics*, the query span “ran in the election” should be linked to the “Candidate” column in the table.

ers compared to the more widely-studied passage-based reading comprehension (RC) problem. Further, TableQA may involve complex questions with multi-cell or aggregate answers.

Most of the TableQA systems use semantic parsing approaches that utilizes language encoders to produce an intermediate logical form from the natural language question which is executed that over the tabular data to get the answer. While some systems (Zhong et al., 2017) were fully supervised, needing pairs of questions and logical forms as training data, more recent systems (Pasupat and Liang, 2015; Krishnamurthy et al., 2017; Dasigi et al., 2019) rely only on the answer as weak supervision and search for a correct logical form. The current best TableQA systems (Herzig et al., 2020; Yin et al., 2020a) capitalize on advances in language modeling, such as BERT, and extend it to encode table representations as well. They are shown to produce excellent results on popular benchmarks such as WikiSQL (Zhong et al., 2017) and WikiTableQuestions(WikiTQ) (Pasupat and Liang, 2015).

With increasing prevalence of text analytics as a centrally-trained service that serves diverse customers, practical QA systems will encounter tables and questions from topics which they may not have necessarily seen during training. It is critical that

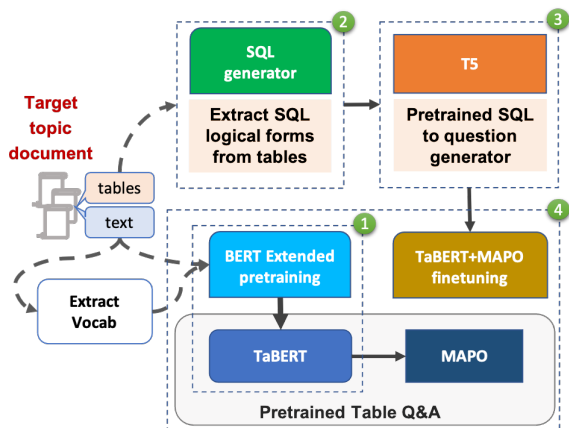


Figure 2: Overview of the proposed T3QA framework for weakly-supervised TableQA.

the language understanding and parsing capabilities of these QA models arising from their training regime are sufficiently robust to answer questions over tables from such unseen topics.

As we show later in this paper, the existing approaches degrade significantly when exposed to questions from topics not seen during training (i.e., topic-shift).² To examine this phenomenon, we first instrument and dissect the performance of these recent systems under topic shift. In particular, we experiment with TaBERT (Yin et al., 2020b), which is a weakly supervised TableQA model which encodes the table and question using BERT-encoder and outputs a logical form using a LSTM decoder. In the example shown in Figure 1, topic shift may cause poor generalization for specific terminology or token usage across unseen topics.

We introduce a novel experimental protocol to highlight the difficulties of topic shift in the context of two well-known Wikipedia-based TableQA datasets: WikiSQL (Zhong et al., 2017) and WikiTableQuestions Pasupat and Liang (2015). Despite recent transformer-based TableQA models being pre-trained with open-domain data, including Wikipedia itself, we observe a performance drop of 5–6% when test instances arise from topics not seen during training.

To address this challenge, we next propose a novel T3QA framework for TableQA training that leads to greater cross-topic robustness. Our approach uses only unlabeled documents with ta-

²Topic shift may be regarded as a case of domain shift studied in the ML community. However, here we refrain from referring to the proposed topic-driven splits as “domains” due to the open-domain nature of these datasets and the pre-training data used to build these models.

bles from the never-seen topic (which we interchangeably call the *target* topic), without any hand-created (question, logical form) pairs in the target topic. Specifically, we first extend the vocabulary of BERT for the new topic. Next, it uses a powerful text-to-text transfer transformer module to generate synthetic questions for the target topic. A pragmatic question generator first samples SQL queries of various types from the target topic table and transcribes them to natural language questions which is then used to finetune the TableQA model on target topic. Finally, T3QA improves the performance of the TableQA model with a post-hoc logical form re-ranker, aided by entity linking. The proposed improvements are applicable to any semantic parsing style TableQA with transformer encoders and is shown to confer generally cumulative improvements in our experiments. To the best of our knowledge, this is the first paper to tackle the TableQA problem in such a zero-shot setting with respect to target topics.

The main contributions of this work are:

- This is the first work to address the phenomenon of *topic shift* in Table Question Answering systems.
- We create novel experimental protocol on 2 existing TableQA datasets to study the effects of topic shift. (WikiSQL-TS and WikiTQ-TS)
- We propose new methods that uses unlabeled text and tables from target topic to create TableQA models which are more robust to topic shift.

2 Related work

Most TableQA systems take a semantic parsing view (Pasupat and Liang, 2015; Zhong et al., 2017; Liang et al., 2017) for question understanding and produce a logical form of the natural language question. Fully-supervised approaches, such as by (Zhong et al., 2017) need pairs of questions and logical form for training. However, obtaining logical form annotations for questions at scale is expensive. A simpler, cheaper alternative is to collect only question-answer pairs as weak supervision (Pasupat and Liang, 2015; Krishnamurthy et al., 2017; Dasigi et al., 2019). Such systems search for the correct logical forms under syntactic and semantic constraints that produce the correct answer. Weak supervision is challenging, owing

Topics	Member sub-topics from Wikipedia	WikiSQL-TS			WikiTQ-TS		
		Train	Dev	Test	Train	Dev	Test
Politics	Crime, Geography, Government, Law, Military, Policy, Politics, Society, World	7728	1236	2314	1836	545	580
Culture	Entertainment, Events, History, Human behavior, Humanities, Life, Culture, Mass media, Music, Organizations	11804	1734	3198	2180	502	691
Sports	Sports	26090	4016	7242	4867	1195	1848
People	People	6548	861	1957	1946	420	743
Misc	Academic disciplines, Business, Concepts, Economy, Education, Energy, Engineering, Food and Drink, Health, Industry, Knowledge, Language, Mathematics, Mind, Objects, Philosophy, Religion, Nature, Science and technology, Universe	3059	395	852	1032	357	438

Table 1: Statistics of the proposed WikiSQL-TS and WikiTQ-TS benchmarks per topic.

to the large search space that includes many possible spurious logical forms (Guu et al., 2017) that may produce the target answer but not an accurate logical transformation of the natural question.

Recent TableQA systems (Herzig et al., 2020; Yin et al., 2020a; Glass et al., 2021) extend BERT to encode the entire table including headers, rows and columns. They aim to learn a table-embedding representation that can capture correlations between question keywords and target cell of the table. TAPAS (Herzig et al., 2020) and RCI (Glass et al., 2021) are designed to answer a question by predicting the correct cells in the table in a truly end-to-end manner. TaBERT (Yin et al., 2020a) is a powerful encoder developed specifically for the TableQA task. TaBERT jointly encodes a natural language question and the table, implicitly creating (i) entity links between question tokens and table-content, and (ii) relationship between table cells, derived from its structure. To generate the structured query, the encoding obtained from TaBERT is coupled with a memory augmented semantic parsing approach (MAPO) (Liang et al., 2018).

Question generation (QG) (Liu et al., 2020; Sultan et al., 2020; Shakeri et al., 2020) has been widely explored in reading comprehension (RC) task to reduce the burden of annotating large volumes of Q-A pairs given a context paragraph. Recently, Puri et al. (2020) used GPT-2 (Radford et al., 2019) to generate synthetic data for RC, showing that synthetic data alone is sufficient to obtain state-of-art on the SQUAD1.1 dataset. For the QG task in TableQA, systems proposed by Benmalek et al. (2019); Guo et al. (2018); Serban et al. (2016) utilize the structure of intermediate logical forms (e.g., SQL) to generate natural language questions. However, none of these QG methods utilize the additional context like table headers, structure and se-

mantics of the tables or the nuances of different possible question types like complex aggregations. To the best of our knowledge, our approach is the first to generate questions specifically for TableQA with the assistance of a logical query and large pre-trained multitask transformers.

Domain adaptation approaches in QA (Lee et al., 2019; Ganin et al., 2016) have so far mostly used adversarial learning with an aim to identify domain agnostic features, including in RC applications (Wang et al., 2019; Cao et al., 2020). However, for the TableQA systems using BERT-style language models with vast pre-training, topic shifts remain an unexplored problem.

3 T3QA framework

To our knowledge, this is the first work to explore TableQA in unseen topic setting. Consequently, no public topic-sliced TableQA dataset is available. We introduce a *topic-shift* benchmark by creating new splits in existing popular TableQA datasets: WikiSQL (Zhong et al., 2017) and WikiTQ (Paspapat and Liang, 2015). The benchmark creation process is described in Section 3.1. Then, we introduce the proposed framework (illustrated in Figure 2) to help TableQA system cope with topic shift. Section 3.2 describes the topic specific vocabulary extension for BERT, followed by Question Generation in target topic in Section 3.3 and reranking logical forms in Section 3.4.

3.1 TableQA topic-shift benchmark

To create a topic-shift TableQA benchmark out of existing datasets, topics have to be assigned to every instance. Once topics are assigned, we create train-test splits with topic shift. I.e., train instances and test instances come from non-overlapping sets of topics. TableQA instances are triplets of the

form {table, question, answer}. For the datasets WikiSQL and WikiTQ, these tables are taken from Wikipedia articles. WikiSQL has 24,241 tables taken from 15,258 articles and WikiTQ has 2,108 tables from 2,104 articles.

The Wikipedia category graph (WCG) is a dense graph organized in a taxonomy-like structure. For the Wikipedia articles corresponding to tables in WikiSQL and WikiTQ, we found that they are connected to 16000+ categories in WCG on an average. Among the *Wikipedia Category:Main topic articles*, Wikipedia articles were connected to 38+ out of 42 categories in WCG.

We use category information from Wikipedia articles to identify topics for each of the article and then transfer those topics to the corresponding tables. The main steps are listed below; details can be found in Appendix B.

- We identify 42 main Wikipedia categories.
- For each table, we locate the Wikipedia article containing it.
- From the page, we follow category ancestor links until we reach one or more main categories.
- In case of multiple candidates, we choose one based on the traversed path length and the number of paths between the candidate and the article.

We cannot take an arbitrary subset of topics for train and the rest for test split to create a topic-shift protocol, because many topics are strongly related to others. For example, topic Entertainment is more strongly related to Music than to Law. To avoid this problem, we cluster these Wikipedia main topics into groups such that similar topics fall in the same group. Using a clustering procedure described in Appendix B, we arrive at 5 high-level topic groups: Politics, Culture, Sports, People and Miscellaneous.

Table 1 gives the membership of each topic group and the number of instances in WikiSQL and WikiTQ dataset per topic. For ease of discussion, we will be calling the five topic groups as topics from now on. For both datasets, we create five leave-one-out topic-shift experiment protocols where in each topic becomes the test set, called the *target topic* and the rest four the training set is called the *source topic(s)*.

In our protocol, for training, apart from the instances from source topic, we also provide tables and document from the target topic. Documents are the text crawled from the target topic articles from Wikipedia. Collecting unlabeled tables and

text data for a target topic is inexpensive. We name these datasets WikiSQL-TS (WikiSQL with topic shift) and WikiTQ-TS.

3.2 Topic specific BERT vocabulary extension

Sub word segmentation in BERT has a potential risk of segmenting named entities or in general unseen words in the target corpus. Vocabulary extension ensures that topic specific words are encoded in entirety and avoids splitting into sub-words. Our goal is to finetune BERT with extended vocabulary on topic specific target corpus to learn topic sensitive contextual representation. So we add frequent topic-specific words to encourage the BERT encoder to learn better topic sensitive representation, which is crucial for better query understanding and query-table entity linking.

3.3 Table-question generation

In our proposed topic-shift experiment protocol with the training set from source topic, unlabeled tables and free text from target topic are provided in the training phase. We propose to use tables from the target topic to generate synthetic question-answer pairs and use these augmented instances for training the TableQA model. Unlike question generation from text, a great deal of additional control is available when generating questions from tables. Similar to Guo et al. (2018), we first sample SQL queries from a given table, and then use a text-to-text transformers (T5) (Raffel et al., 2020) based sequence-to-sequence model to transcribe the SQL query to a natural language question.

3.3.1 SQL sampling

For generating synthetic SQL queries from a given table T, we have designed a focused and controllable SQL query generation mechanism presented in Algorithm 1. Our approach is similar to Zhong et al. (2017) but unlike the existing approaches, we use guidance from target query syntax to offer much more control over the type of natural language questions being generated. We also use additional context such as table header, target answer cell to help the model generate more meaningful questions suitable for T3QA. We sample the query type (simple retrieval vs. aggregations) and associated where clauses from a distribution that matches the prior probability distribution of training data, if that is available. Sampling of query type and number of where clauses is important to mitigate

Algorithm 1: Algorithm for SQL query Generation from a table.

```

Input: Table T, Integer targetNum
Output: SQLQuery[] generatedSQLs
1 foreach Column C ∈ T.columns do
2   | C.dataType ← ExtractDataType(C)
3 while generatedSQLs.size() < targetNum do
4   SQLQuery S ← Empty ;
5   Integer num_where ← sample from a multinomial
   distribution over {1,2,3,4} ;
6   Where[] whereClauses ← generateWhereClauses(T,
   num_where) ;
7   S.Where ← whereClauses ;
8   String returnType ← sample from a multinomial distribution
   over {SELECT, SUM, AVG, MAX, MIN} ;
9   if returnType ∈ {SUM, AVG, MAX, MIN} then
10  | Column selectColumn ← sample Column C from
   {T.Columns \ whereClauses.Columns |
   C.dataType = Numeric} ;
11 else
12  | Column selectColumn ← sample Column C from
   {T.Columns \ whereClauses.Columns}
13  S.Select ← selectColumn ;
14  if ∃ SQL S2 | S2.Where ⊂ S.Where AND
   S2.Result = S1.result then
15  | valid = false;
16  if returnType ∈ {SUM, AVG, MAX, MIN} AND
   numOfRows(S.Result) == 1 then
17  | valid = false;
18  if valid then
19  | generatedSQLs.add(S)
20 Return generatedSQLs

```

the risk of learning a biased model that cannot generalize for more complex queries with more than 2 where clauses, as reported by Guo et al. (2018).

The generated SQL queries are checked for various aspects of semantic quality, beyond mere syntactic correctness in typical rule based generations. WikiSQL has a known imitation: even an incorrect SQL query can produce the same answer as the gold SQL query. To avoid such cases, we make two important checks: (1) The WHERE clauses in the generated SQL queries must all be mandatory to produce the correct answer. i.e., dropping a WHERE clause should not produce the expected answer and (2) a generated SQL query with an aggregation must have at least 2 rows to aggregate on and therefore, dropping the aggregation will not produce the expected answer. These quality checks ensure that the generated synthetic SQL queries are fit to be used in TableQA training pipeline.

3.3.2 T5 transfer learning for QG

For question generation in the TableQA setup, it is more intuitive to create SQL queries first and then use the structure of the SQL query to translate it to a natural language question. Previously, Guo et al. (2018) and Benmalek et al. (2019) used LSTM-based sequence to sequence models for direct question generation from tables. However, we

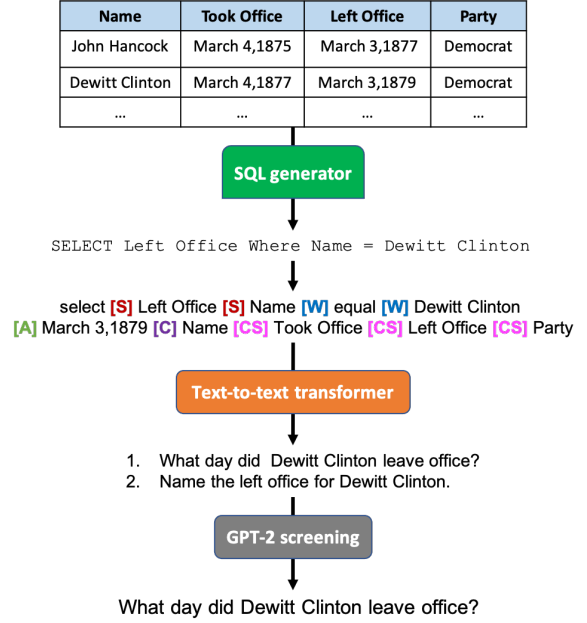


Figure 3: Generating synthetic questions on target topics using only tables. Special tokens are shown in colored font.

hypothesize that apart from SQL queries, using answers and column headers with the help of transformer based models, can be more effective.

For our question generation module we have used unified text-to-text transformers (T5) (Raffel et al., 2020), which is popular for its *constrained* text generation capabilities for multiple tasks such as translation and summarization. To leverage this capability of T5 for generating natural language questions from SQL queries, we encode a SQL query in a specific text format. We also pass the answer of the SQL query and the column headers of table to T5 as we observe that using these two sets of extra information along with the SQL query helps in generating better questions, especially with "Wh" words. As illustrated in Figure 3, the generated SQL query with answer and column headers are encoded into a specific sequence before passing onto T5 model. Special separator tokens are used to demarcate different parts of the input sequence: [S] to specify the main column and operation, [W] demarcates elements in a WHERE clause, [A] marks the answer, [C] and [CS] show the beginning of set of column headers and separation between them, respectively.

In this example, one can observe that although the SQL query do not have any term on day or date, our QG module was able to add "What day". Furthermore, ill-formed and unnatural questions gener-

Type	Ground truth SQL	Generated Question	Ground truth question
Lookup	SELECT Rounds WHERE Chassis = b195	What round has a car with a b195 chassis?	Which rounds had the B195 chassis?
	SELECT College WHERE Player = Paul Seiler	What college does Paul Seiler play for?	What college has Paul Seiler as a player?
	SELECT Date WHERE Attendance > 20,066 AND Home = Tampa Bay	On what date was the attendance more than 20,066 at Tampa Bay?	When has an Attendance larger than 20,066 in tampa bay?
Aggregate	SELECT SUM(Attendance) WHERE Date = May 31	How many people attended the May 31 game?	How many people attended the game on May 31?
	SELECT MAX(Mpix) WHERE Aspect Ratio = 2:1 AND Height < 1536 AND Width < 2048	What is the highest Mpix with an Aspect Ratio of 2:1, a Height smaller than 1536, and a Width smaller than 2048?	What camera has the highest Mpix with an aspect ratio of 2:1, a height less than 1536, and a width smaller than 2048?
	SELECT AVG(Score) WHERE Player = Lee Westwood	What is Lee Westwood's average score?	What is the average score with lee westwood as the player?

Table 2: Ground truth SQL queries with generated questions (using T5 based QG module) and gold questions

Operation	Sampled SQL	Generated Question
SELECT	SELECT Production code WHERE Written by = José Rivera	what is the production code for the episode written by José rivera?
	SELECT Average WHERE Rank by average > 3 AND Number of dances=17	what is the average for a rank by average larger than 3 and 17 dances?
MAX	SELECT MAX(SEATS) WHERE Kit/Factory = Factory	can you tell me the highest seats that has the kit/factory of factory?
	SELECT MAX(YEAR) WHERE WINS = 70 AND Manager = Jim Beauchamp	what is the most recent year of the team with 70 wins and manager Jim Beauchamp?
MIN	SELECT MIN(Rank) WHERE Nationality = RUS	which rank is the lowest one that has a nationality ofrus?
	SELECT MIN(Televote Points) WHERE Panel Points = 0	which Televote points is the lowest one that has panels pointss of 0?
SUM	SELECT SUM(Game) WHERE Team = Baltimore	what is the sum of game, when team is Baltimore?
	SELECT SUM(Division) WHERE Year < 2011 AND Playoffs = Did not qualify	what is the total number of division(s), when year is less than 2011, and when playoffs did not qualify?
AVG	SELECT AVG(Digital PSIP) WHERE Network = Omni Television	which digital PSIP has a network of Omni television?
	SELECT AVG(Attendance) WHERE Week < 5	what was the average attendance before week 5?

Table 3: Synthetic questions generated on *sampled SQLs* with SELECT and various aggregate functions on WikiSQL-TS tables. Observe that the quality of questions is generally better with SELECT operation than aggregate ones. The reason for this might be that the data used to train QG module includes more SELECT questions.

ated by T5 model are filtered out using a pretrained GPT-2 model (Radford et al., 2019). We removed questions with the highest perplexity scores before passing the rest to the TableQA training module.

For training the QG module, we use SQL queries and questions provided with the WikiSQL dataset. In our experiments, only query+question pairs from the source topics are used to train the question generation module and synthetic questions are gener-

ated for the target topic.

We are able to produce high-quality questions using this T5 transcription. Table 2 shows a few example of generated questions from ground truth SQL and Table 3 on sampled SQLs. Observe that the model is able to generate lookup questions, multiple conditions, and aggregate questions of high quality. It is interesting to see that for the first example in Table 2, T5 model included the

term *car* in the question even though it was not available in the SQL query, probably taking the clue from *chassis*. Some questions created from sampled SQLs for WikiTQ tables is provided in Appendix C.

3.4 Reranking logical forms

We analysed the logical forms predicted by TaBERT model in WikiSQL-TS and observed that the top logical forms often do not have the correct column headers and cell values. In fact, in WikiSQL-TS there is a 15–20% greater chance of finding a correct prediction from the top-5 predicted logical form than the top 1.

We propose to use a classifier, Gboost (Friedman, 2002) to rerank the predicted top-5 logical form. Given a logical form and table-question pair we create a set of features on which a classifier is trained to give higher score to the correct logical form.

The logical form-question pair which gives the correct prediction is labelled as *+ve* and wrong predictions as *-ve*. We use the predicted logical forms for source topic dev set to train this classifier and in the inference step while predicting for target topic, the logical form which got highest score by the classifier is selected.

3.4.1 Features for logical form reranker

Two sets of features are extracted for the reranker: (1) entity linking based features, (2) logical form based features.

Entity linking based features: This captures matches between query fragments and table elements. Our system of entity linking using string matching also finds partial matches. Partial matches happen when only a part of column name or cell value appear in the question. Another scenario is when token in the question partially matches with multiple entities in the table. We create three feature separately for cell values and column headers.

- Number of linked entities in logical form which appear partially or fully in question.
- Sum of ratio of tokens matched with entities in logical form. If the questions has word *States* and corresponding entity in table is *United States*, then the ratio would be 0.5.
- Sum of a measure of certainty in entity linking. if the question token partially matches with multiple entities in table then certainty is less. If the question has word *United* and there are three en-

ties in the table *United Kingdom, United States* and *United Arab Emirates*, then we assign certainty score as 1/3.

Only logical form features:

- Probability score of logical form given by the TableQA model
- Length of answer obtained by using this logical form. Length here doesn't mean the number of characters but number of cells in prediction.
- If 'count' is present in the logical form
- If 'select' is present in the logical form
- Number of where clauses.
- If columns are repeated in the logical form.

4 Experiments and Analysis

Here we describe key details of the experimental setup, the models compared and evaluation techniques. We also provide a thorough analysis of the results to highlight the key takeaways.

4.1 Setup

We consider WikiSQL-TS and WikiTQ-TS for our experiments with topic assignments as described in Section 3.1. The larger WikiSQL-TS dataset consists of tables, questions and corresponding ground truth SQL queries, whereas WikiTQ-TS contains only natural language questions and answers. The five topics are 1) *Politics* 2) *Culture* 3) *Sports* 4) *People* and 5) *Miscellaneous*. Table 1 captures some interesting statistics about the topic split benchmark created from WikiSQL. All experiments are conducted in a leave-one-out (LOO) fashion where the target topic examples are withheld. For example, if the target topic is *Politics* then the model is trained using the train set and dev set of *Culture, Sports, People, Misc* and evaluated on test set of *Politics*. Further, a **composite dev set** is curated by adding equal number of synthetically generated questions from the target topic to the dev set of source topics.

4.2 Models

We perform all experiments using a variant of TaBERT+MAPO³ architecture, with the underlying BERT model initialized with `bert-base-uncased`. **TaBERT+MAPO** uses standard BERT as table-question encoder and MAPO (Liang et al., 2018) as the base semantic parser. **TaBERT_t+MAPO** uses topic spe-

³https://github.com/pcyin/pytorch_neural_symbolic_machines

Topic	TaBERT	TaBERT _t	TaBERT+QG	TaBERT _t +QG	TaBERT _t +QG+Reranker
Politics	61.71	64.95	64.26	66.12	70.22
Culture	64.89	66.10	69.32	69.88	72.63
Sports	62.10	62.70	63.03	63.83	66.5
People	60.34	61.93	63.10	66.27	70.87
Misc	61.85	59.03	64.31	64.43	69.60

Table 4: Performance on WikiSQL-TS benchmark. Here, TaBERT means TaBERT+MAPO and TaBERT_t means TaBERT_t+MAPO. All numbers are in %.

cific pre-trained BERT encoder (as described in section 3.2). Similar to the base model, this model use MAPO as the base semantic parser. **TaBERT+MAPO+QG** uses an extended training set with question answer pairs generated from the proposed QG model to train the TaBERT+MAPO model. **TaBERT_t+MAPO+QG** uses an initialized BERT encoder parameters with topic specific pre-trained BERT and add question-answer pairs generated by our QG model to train the TaBERT_t+MAPO model.

Table Question Generation (QG): We use the T5 implementation of Wolf et al. (2019) for question generation, initialized with `t5-base` and finetuned using SQL and corresponding questions from WikiSQL dataset. To ensure that the target topic is not leaked through the T5 model, we trained five *topic-specific* T5 models, one for each leave-one-out group by considering only SQL-question pairs from the source topic only. As WikiTQ-TS does not have ground truth SQL queries included in the dataset, we use T5 trained on WikiSQL-TS to generate synthetic questions. We use a batch-size of 10 with a learning rate of 10^{-3} .

Implementation details: We build upon the existing code base for TaBERT+MAPO released by Yin et al. (2020b) and use BERT_{base} as the encoder for tables and questions. We use topic-specific vocabulary (explained in Section 3.2) for BERT’s tokenizer and train it using MLM (masked language model) objective for 3 epochs with $p=0.15$ chance of masking a topic-specific high frequency (occurring more than 15 times in target topic corpus) token. We optimize BERT parameters using Adam optimizer with learning rate of 5×10^{-5} .

All numbers reported are from the test fold, fixing system parameters and model selection with best performance on the corresponding composite dev set. Further details and the dataset are provided in the supplementary material.

Topic	Number of WHERE clauses			
	1	2	3	4
Politics	2.11	12.24	6.66	15.00
Culture	0.85	8.93	4.89	5.00
Sports	0.96	6.81	5.20	-3.89
People	1.65	9.52	11.03	6.25
Misc	1.71	13.00	10.00	33.34

Table 5: Change in performance in WikiSQL-TS after applying Reranker to TaBERT_t+MAPO+QG, across number of WHERE clauses. All numbers are in absolute %.

4.3 Results and Analysis

WikiSQL-TS: TaBERT_t+MAPO improves over TaBERT+MAPO for four out of five test topics by an average of 1.66%, showing the advantage of vocabulary extension (Table 4). In addition to supplying the topic-specific *sense* of vocabulary, fine tuning also avoids introducing word-pieces that adversely affect topic-specific language understanding. For instance, for the topic *culture* the whole word ‘rockstar’ is added to the vocabulary rather than the word-pieces ‘rocks’, ‘##tar’. We implement vocabulary extension by using the 1000 placeholders in BERT’s vocabulary, accommodating high frequency words from the target topic corpus.

Further, TaBERT+MAPO+QG significantly outperforms TaBERT+MAPO and also TaBERT_t+MAPO when finetuned with target topic samples obtained from QG (after careful filtering). In WikiSQL-TS, QG also improves the performance of TaBERT_t+MAPO, though relevant vocabulary was already added to BERT, suggesting additional benefits of QG in T3QA framework. While vocabulary extension ensures topical tokens are encoded, QG improves implicit linking between question and table header tokens within the joint encoding of question-table. The largest improvement of 10.53% and 7.74% is obtained for People and Culture respectively. Moreover, TaBERT+MAPO+QG out-performs an in-topic performance of 64.07% and 67% with 66.27% and 69.88% (details in Appendix D), showing that the unseen topic performance can be substantially improved with only auxiliary text and tables from documents without explicitly annotated table, question, and answer tuples.

As mentioned, *Misc* is a topic chimera with a mixed individual statistics, hence an explicit injection of frequent vocabulary does not significantly improve TaBERT_t+MAPO over

Topic	TaBERT+MAPO							TaBERT _t +MAPO+QG+Reranker						
	overall	select	count	min	max	sum	avg	overall	select	count	min	max	sum	avg
Politics	61.71	62.82	66.17	53.28	58.64	46.26	60.21	70.22	73.90	60.59	70.98	56.57	56.71	65.59
Culture	64.89	64.47	70.62	62.74	65.56	62.66	60.71	72.63	74.50	65.01	69.53	69.93	64.0	63.09
Sports	62.10	61.60	57.16	69.55	72.09	54.14	62.07	66.5	67.06	45.45	78.85	74.39	67.15	69.41
People	60.34	59.10	66.92	60.71	69.56	50.72	73.33	70.87	72.55	60.0	72.82	65.17	60.86	73.33
Misc	61.85	60.8	65.0	72.34	76.19	44.82	55.17	69.60	69.76	66.25	95.23 ↑	74.46	44.82	51.72

Table 6: Performance on WikiSQL Topic specific benchmark across various question types. The largest group, *select*, is shown in **bold**. Largest improvement is shown as ↑. All numbers are in absolute %.

Topic	TaBERT	TaBERT _t	TaBERT+QG	TaBERT _t +QG	TaBERT _t +QG+Reranker
Politics	40.52	41.03	41.55	41.38	43.79
Culture	36.03	38.49	38.49	37.05	39.50
Sports	37.55	37.5	37.93	39.12	41.50
People	35.94	37.69	37.42	36.61	39.30
Misc	38.58	40.64	41.10	40.18	42.23

Table 7: Performance on WikiTQ-TS benchmark. Here, TaBERT means TaBERT+MAPO and TaBERT_t means TaBERT_t+MAPO. All numbers are in %.

TaBERT+MAPO. However, TaBERT+MAPO+QG outperforms TaBERT_t+MAPO by 5.4% due to QG, suggesting that the improvement from both methods are disjoint. Further, Question generation, though conditioned on the table and topic specific text is not supplied with the topic vocabulary. We also observe that the composite dev set with 50% real questions and 50% questions generated on tables from target topic improves performance. Tables 4 & 5 take the advantage of ground truth SQL queries to further dissect the performance along question types and number of WHERE clauses.

Number of Where clauses: As described previously, performance of TaBERT+MAPO is substantially affected by the number of WHERE clauses in the ground truth logical form (also observed by (Guo et al., 2018)), see Appendix A. Table 5, shows that performance improvement by “Reranker” is significantly higher for more than 1 WHERE clause. This might have happened because TaBERT+MAPO prefers to decode shorter logical forms, whereas the reranker prioritizes logical forms with more linked entities present from the question.

WikiSQL question types: Table 4 breaks down the performance of TaBERT+MAPO+QG based on the question types labels obtained from the dataset ground truth only for analysis. The improvement, viewed from the lens of question types is more significant with average gain in SELECT-style queries at 9.76%. Aggregate (count, min/max, sum, avg) questions are more challenging to gener-

ate as the answer is not present in the table. Consequently, the performance improvement with QG is less significant for these question types.

WikiTQ-TS: WikiTQ-TS is a smaller dataset and contains more complex questions (negatives, implicit nested query) compared to WikiSQL-TS. Correspondingly, there is also less topic specific text to pretrain the TaBERT encoder. Despite these limitations, we observe in Table 7 that TaBERT_t with vocabulary extension and pretraining shows overall improvement. We resort to using synthetic questions generated from QG model of WikiSQL-TS, due to unavailability of ground truth SQL queries in WikiTQ. Hence, the generated questions are often different in structure from the ground truth questions. Samples of real and generated questions are in Table 8 of Appendix C. Despite this difference in question distribution we see TaBERT+QG consistently performs better than the baseline. We provide an analysis of the perplexity scores from TaBERT and TaBERT_t on the generated questions in Appendix G. Ultimately, the proposed T3QA framework significantly improves performance in all target domains.

5 Conclusion

This paper introduces the problem of TableQA for unseen topics. We propose novel topic split benchmarks over WikiSQL and WikiTQ and highlight the drop in performance of TaBERT+MAPO, even when TaBERT is pretrained on a large open domain corpora. We show that significant gains in performance can be achieved by (i) extending the vocabulary of BERT with topic-specific tokens (ii) fine-tuning the model with our proposed constrained question generation which transcribes SQL into natural language, (iii) re-ranking logical forms based on features associated with entity linking and logical form structure. We believe that the proposed benchmark can be used by the community for building and evaluating robust TableQA models for practical settings.

References

- Ryan Benmalek, Madian Khabsa, Suma Desu, Claire Cardie, and Michele Banko. 2019. Keeping notes: Conditional natural language generation with a scratchpad encoder. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4157–4167.
- Yu Cao, Meng Fang, Baosheng Yu, and Joey Tianyi Zhou. 2020. [Unsupervised domain adaptation on reading comprehension](#).
- Pradeep Dasigi, Matt Gardner, Shikhar Murty, Luke Zettlemoyer, and Eduard Hovy. 2019. [Iterative search for weakly supervised semantic parsing](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2669–2680, Minneapolis, Minnesota. Association for Computational Linguistics.
- Inderjit S Dhillon. 2001. Co-clustering documents and words using bipartite spectral graph partitioning. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 269–274. ACM.
- Jerome H Friedman. 2002. Stochastic gradient boosting. *Computational statistics & data analysis*, 38(4):367–378.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. [Domain-adversarial training of neural networks](#).
- Michael Glass, Mustafa Canim, Alfio Gliozzo, Saneem Chemmengath, Vishwajeet Kumar, Rishav Chakravarti, Avirup Sil, Feifei Pan, Samarth Bhargava, and Nicolas Rodolfo Fauceglia. 2021. Capturing row and column semantics in transformer based question answering over tables. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1212–1224.
- Daya Guo, Yibo Sun, Duyu Tang, Nan Duan, Jian Yin, Hong Chi, James Cao, Peng Chen, and Ming Zhou. 2018. Question generation from sql queries improves neural semantic parsing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1597–1607.
- Kelvin Guu, Panupong Pasupat, Evan Zheran Liu, and Percy Liang. 2017. [From language to programs: Bridging reinforcement learning and maximum marginal likelihood](#).
- Jonathan Herzig, Paweł Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Martin Eisenschlos. 2020. [Tapas: Weakly supervised table parsing via pre-training](#).
- Jayant Krishnamurthy, Pradeep Dasigi, and Matt Gardner. 2017. [Neural semantic parsing with type constraints for semi-structured tables](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1516–1526, Copenhagen, Denmark. Association for Computational Linguistics.
- Seanie Lee, Donggyu Kim, and Jangwon Park. 2019. [Domain-agnostic question-answering with adversarial training](#).
- Chen Liang, Jonathan Berant, Quoc Le, Kenneth Forbus, and Ni Lao. 2017. 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.
- Chen Liang, Mohammad Norouzi, Jonathan Berant, Quoc V Le, and Ni Lao. 2018. Memory augmented policy optimization for program synthesis and semantic parsing. In *Advances in Neural Information Processing Systems*, pages 9994–10006.
- Bang Liu, Haojie Wei, Di Niu, Haolan Chen, and Yancheng He. 2020. [Asking questions the human way: Scalable question-answer generation from text corpus](#). *Proceedings of The Web Conference 2020*.
- Panupong Pasupat and Percy Liang. 2015. [Compositional semantic parsing on semi-structured tables](#).
- Raul Puri, Ryan Spring, Mostofa Patwary, Mohammad Shoeybi, and Bryan Catanzaro. 2020. [Training question answering models from synthetic data](#).
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#).
- Iulian Vlad Serban, Alberto Garcia-Duran, Caglar Gulcehre, Sungjin Ahn, Sarath Chandar, Aaron Courville, and Yoshua Bengio. 2016. Generating factoid questions with recurrent neural networks: The 30m factoid question-answer corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 588–598.
- Siamak Shakeri, Cicero Nogueira dos Santos, Henry Zhu, Patrick Ng, Feng Nan, Zhiguo Wang, Ramesh Nallapati, and Bing Xiang. 2020. [End-to-end synthetic data generation for domain adaptation of question answering systems](#).
- Md Arafat Sultan, Shubham Chandel, Ramón Fernández Astudillo, and Vittorio Castelli. 2020. [On the importance of diversity in question generation for](#)

- QA. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5651–5656, Online. Association for Computational Linguistics.
- Huazheng Wang, Zhe Gan, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, and Hongning Wang. 2019. [Adversarial domain adaptation for machine reading comprehension](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2510–2520, Hong Kong, China. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, pages arXiv–1910.
- Pengcheng Yin, Graham Neubig, Wen tau Yih, and Sebastian Riedel. 2020a. [Tabert: Pretraining for joint understanding of textual and tabular data](#).
- Pengcheng Yin, Graham Neubig, Wen tau Yih, and Sebastian Riedel. 2020b. [TaBERT: Pretraining for joint understanding of textual and tabular data](#). In *Annual Conference of the Association for Computational Linguistics (ACL)*.
- Torsten Zesch and Iryna Gurevych. 2007. Analysis of the wikipedia category graph for nlp applications. In *Proceedings of the Second Workshop on TextGraphs: Graph-Based Algorithms for Natural Language Processing*, pages 1–8.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. [Seq2sql: Generating structured queries from natural language using reinforcement learning](#).

Topic Transferable Table Question Answering (Appendix)

A TaBERT performance on WikiSQL

We analyse accuracy of TaBERT model on WikiSQL in terms of the number of WHERE clauses, which are skewed as shown in Fig. 4(a). In Fig. 4(b), we observe that accuracy decreases when ground truth SQL has a larger number of WHERE clauses. Interestingly, we observe in Fig. 4(c) and (d) that even though the model achieves 30% to 40% accuracy for 2–4 WHERE clauses, the predicted logical form still produced one WHERE clause. This shows that, for many questions, wrong or incomplete logical forms can produce correct answers.

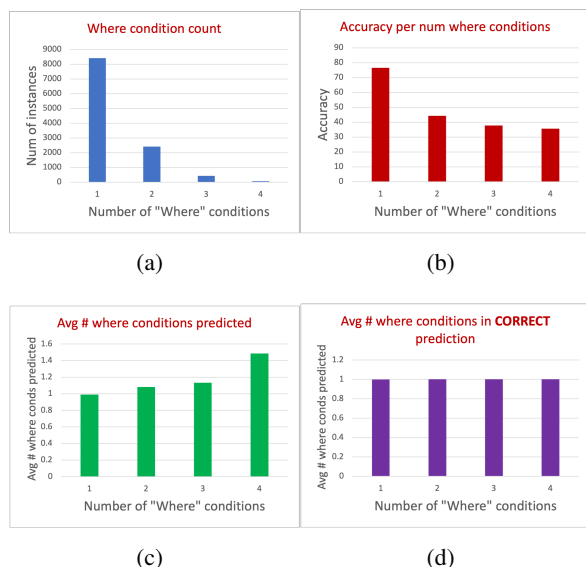


Figure 4: WHERE clause analysis on TaBERT+MAPO performance on SELECT questions in WikiSQL test set (not topic shift): (a) Frequency of questions in different number of WHERE clauses buckets; (b) Accuracy achieved in each WHERE bucket; (c) Average number of conditions in predicted logical forms; (d) Average number of conditions in predicted logical forms which produces correct answers.

B Topic-shift benchmark details

Continuing from Section 3.1, this section provides more details about the creation of the topic shift benchmark datasets. Each Wikipedia article is tagged with a set of categories and each category is further tagged with a set of parent categories, and those to their parent categories, and so on. The whole set of Wikipedia categories are organized in

a taxonomy-like structure called Wikipedia Category Graph (WCG) (Zesch and Gurevych, 2007). These categories range from specific topics such as "Major League Soccer awards" to general topics such as "Human Nature". To have categories of similar granularity, we use the 42 categories listed in *Wikipedia Category:Main topic articles*⁴ as topics.

To assign a unique category to a Wikipedia article, we proceed as follows:

- For each Table T , we extract the Wikipedia Article A which contains Table T .
- We start with the category of A and traverse the hierarchical categories till we reach one (or more) of the 42 categories listed in *Wikipedia Category:Main topic articles*.
- If multiple main topic categories can be reached from A , we take the category which is reached via the shortest path (in terms of number of hierarchical categories traversed from A) and assign that as the category for table T .
- If there are multiple main topic categories which can be reached with the same length of shortest path, we consider the number of different paths between the main topic category and A as the tie breaker to assign the topic for A .

Now we describe the method used to cluster categories into topics. For every article we identify five categories closest to the article in Wikipedia Category Graph. We then compute the Jaccard similarity between two topics as the ratio of number of common articles between topics (in the first-5 list) to the total number of articles assigned to both topic. Using this similarity, we apply spectral co-clustering (Dhillon, 2001) to form five topic groups.

To verify the coherence of the five topic groups, we performed a vocab overlap exercise. For questions in WikiTQ, we find the 100 most frequent words in the test set of each of the topics. Then we measure how many of these frequent words appeared in the train set of each of these topics. Table 9 shows that word overlap is large within clusters.

⁴https://en.wikipedia.org/wiki/Category:Main_topic_articles

Gold questions in the dataset	Generated Questions
<ul style="list-style-type: none"> - how many v8 engines competed in the 1982 British formula one season? - how many entrants have names that contain the word "team"? - name an entrant with no cosworth engines. - how many drivers use v8 engines? - what is the total number of drivers listed? - who is the only driver to use a v12 engine? - Are there any other engines listed besides cosworth or brm? - Which is the only driver whose vehicle used a brm 202 v12 engine? - What is the last chassis listed? 	<ul style="list-style-type: none"> - which constructor has an Entrant of Colin Bennett racing, and a no smaller than 7.0? - what is the average number that has a constructor of Ensign? - who is the driver with a no of 14.0? - what is the average number of FW07 chassisassis? - what engine has a driver of Steveo'rouke? - name the most number for Chassis being N180b. - what is the lowest number that has a constructor of Ensign? - what is the total number that has a constructor of Williams? - what is the largest number for teamensign?

Table 8: Real questions and generated questions for table `csv/203-csv/1.tsv` from WikiQ-TS dataset. Observe that generated questions (finetuned on WikiSQL-TS) have different semantics than the real questions.

Test/Train	Politics	Culture	Sports	People	Misc
Politics	88	73	74	72	64
Culture	79	87	89	85	69
Sports	67	72	100	81	50
People	66	78	88	93	55
Misc	74	73	74	72	68

Table 9: Percent vocabulary match within and across topics (category groups/clusters).

C Questions generation for target topics

Table 8 compares ground truth questions with that of generated questions for the same table from WikiSQL-TS. One can see that even template of questions in real dataset is very different and often tougher than the generated ones. Question generator being trained on WikiSQL-TS dataset with much simple questions might be the reason for this phenomenon.

D Performance when topics are seen

We further analyse the performance of the model in both seen-topic training (when the topic specific train set is available), against the unseen topic train (when the topic specific train set is not used during training). In Table 10, we present results in both training setups.

Topic	Seen Topic	Unseen Topic
Politics	65.52	61.71
Culture	67.26	64.88
Sports	63.14	62.10
People	64.07	60.34
Misc	63.14	61.85

Table 10: Drop in performance due to topic shift in WikiSQL-TS. (Numbers are percentages.)

E Additional Experiments

Table 11 shows the absolute values corresponding to Table 6. in the paper. The performance of both models is lower for questions with larger WHERE clauses. Table 12 summarizes the answer accuracy of TaBERT_t+MAPO +QG +Reranker and TaBERT+MAPO across number of where clauses in the ground truth logical forms.

Topic	Number of WHERE clauses			
	1	2	3	4
Politics	75.78/73.67	58.36/46.12	51.66/45.0	40.0/25.0
Culture	77.52/76.67	61.23/52.30	52.44/47.55	55.0/50.0
Sports	71.26/70.30	57.62/50.81	56.26/51.06	48.05/51.94 ↓
People	77.61/75.96	61.37/51.85	58.82/47.79	25.0/18.75
Misc	75.17/73.46	57.0/44.0	54.0/44.0	66.67/33.33

Table 11: WikiSQL-TS performance for TaBERT_t+MAPO +QG+Reranker and TaBERT_t+MAPO+QG (seperated by '/') across number of WHERE clauses in the ground truth logical forms. All numbers are in %.

Topic	Number of WHERE clauses			
	1	2	3	4
Politics	75.78/67.61	58.36/46.94	51.66/48.33	40.0/40.0
Culture	77.52/71.36	61.23/46.46	52.44/51.05	55.0/35.0
Sports	71.26/67.78	57.62/50.66	56.26/53.9	48.05/44.16
People	77.61/69.55	61.37/44.62	58.82/45.59	25.0/37.5 ↓
Misc	75.17/70.58	57.0/40.5	54.0/46.0	66.67/66.67

Table 12: WikiSQL-TS performance for TaBERT_t+MAPO +QG +Reranker and TaBERT+MAPO (seperated by '/') across number of WHERE clauses in the ground truth logical forms.

F Training details

We train all TaBERT+MAPO variants for 10 epochs on 4 Tesla V100 GPUs using mixed precision training⁵. For training TaBERT+MAPO, we set batch size to 10, number of explore samples 10 and other hyperparameters are kept same as (Yin et al., 2020a). We build upon codebase⁶ released by (Yin et al., 2020a). The hyper-parameters (where not mentioned explicitly) are the same as the original code. We include all the data splits and predictions from our best model as supplementary material with the paper. These will be released publicly upon acceptance. The experimentation requires for 5 topics, we performed 6 variations of the model. We performed search over 4 sets of hyper-parameters, primarily on the composition of generated vs. real questions.

G TaBERT vs. TaBERT_t perplexity of generated questions for WikiTQ-TS

We compute the perplexity scores over a subset of 50 generated questions used in the experiments using both TaBERT and TaBERT_t language models. Note that TaBERT is pretrained on large open domain set whereas TaBERT_t was further fine-tuned on topic specific documents closely related to the tables of target domain. As shown in Table 13, the average perplexity score from TaBERT_t is larger than TaBERT. This indicates that the generated questions are not aligned to the topic in the case of WikiTQ-TS. This is due to the lack of any training examples for specific to the dataset, as mentioned in Section 4.3. Future work on topic-specific question generation may address this issue.

Topic	TaBERT	TaBERT _t
Politics	1.088	1.112
Culture	1.099	1.142
Sports	1.084	1.134
People	1.109	1.164
Misc	1.104	1.153

Table 13: The average perplexity scores of a subset of generated questions from TaBERT and TaBERT_t for WikiTQ-TS

We suspect that this might be the reason why TaBERT_t+QG does not outperform TaBERT+QG

⁵<https://github.com/NVIDIA/apex>

⁶https://github.com/pcyin/pytorch_neural_symbolic_machines

in the case of WikiTQ-TS (Table 7). However, we obtain best performance via the overall T3QA framework.