

Localize, Retrieve and Fuse: A Generalized Framework for Free-Form Question Answering over Tables

Wenting Zhao¹ Ye Liu² Yao Wan³ Yibo Wang¹ Zhongfen Deng¹ Philip S. Yu¹

¹Department of Computer Science, University of Illinois at Chicago, IL, USA

²Salesforce Research, Palo Alto, USA

³School of Computer Sci. & Tech., Huazhong University of Science and Technology, China

{wzhao41, ywang633, zdeng21, psyu}@uic.edu

yeliu@salesforce.com, wanyao@hust.edu.cn

Abstract

Question answering on tabular data (*a.k.a* TableQA), which aims at generating answers to questions grounded on a provided table, has gained significant attention recently. Prior work primarily produces concise factual responses through information extraction from individual or limited table cells, lacking the ability to reason across diverse table cells. Yet, the realm of free-form TableQA, which demands intricate strategies for selecting relevant table cells and the sophisticated integration and inference of discrete data fragments, remains mostly unexplored. To this end, this paper proposes a generalized three-stage approach: *Table-to-Graph conversion and cell localizing*, *external knowledge retrieval*, and *the fusion of table and text* (called TAG-QA), to address the challenge of inferring long free-form answers in generative TableQA. In particular, TAG-QA (1) locates relevant table cells using a graph neural network to gather intersecting cells between relevant rows and columns, (2) leverages external knowledge from Wikipedia, and (3) generates answers by integrating both tabular data and natural linguistic information. Experiments showcase the superior capabilities of TAG-QA in generating sentences that are both faithful and coherent, particularly when compared to several state-of-the-art baselines. Notably, TAG-QA surpasses the robust pipeline-based baseline TAPAS by 17% and 14% in terms of BLEU-4 and PARENT F-score, respectively. Furthermore, TAG-QA outperforms the end-to-end model T5 by 16% and 12% on BLEU-4 and PARENT F-score, respectively.¹

1 Introduction

Question answering is to generate precise answers by interacting efficiently with unstructured, structured, or heterogeneous contexts, such as paragraphs, knowledge bases, tables, images, and various combinations thereof (Burke et al., 1997; Yao

¹Source code will be released at <https://github.com/wentinghome/TAGQA>.

Rank	Rider	Team	Speed	Time
1	Northern Ireland Robert D	Yamaha	102.46 mph	1:28.22.2
2	Scotland Steve Hislop	Yamaha	101.27	1:29.24.8
3	Wales Ian Lougher	Yamaha	100.62	1:29.59.2

[Q]: Who won in the first three places of The Newcomers Manx Grand Prix race?
[T5]: Scotland's [Steve Hislop](#) finished second with 101.27 mph and [Wales' Ian Lougher](#) finished third with 100.62 mph.
[TAPAS]: The Newcomers Manx Grand Prix race was won by the [Spaniard](#) in the first three places with a time of 1:28.22.2 seconds, 1:29.24.8 seconds and a time of 1:29.59.2.
[MATE]: [Robert Dunlop](#) won in the first three places, followed by [Steve Hislop](#) in the second.
[Ours]: The Newcomers Manx Grand Prix race was won by [Robert Dunlop](#) from [Scotland](#) [Steve Hislop](#) in 2nd place and [Ian Lougher](#) in 3rd place at 100.62 mph.
[Reference]: The Newcomers Manx Grand Prix race was won by [Robert Dunlop](#) from [Steve Hislop](#) in 2nd place and [Ian Lougher](#) in 3rd place at a race speed of 100.62.

Figure 1: A motivating example to show the insights of our proposed approach when comparing with several state-of-the-art methods.

and Van Durme, 2014; Talmor et al., 2021; Hao et al., 2017). Among these, question answering on tabular data (TableQA) is a challenging task that requires the understanding of table semantics, as well as the ability to reason and infer over relevant table cells (Herzig et al., 2021; Chen et al., 2020b, 2021b).

For the task of TableQA, from our investigation, most current studies are focusing on the factoid TableQA, in which the answer is in a few words or a phrase copied directly from relevant table cells. In particular, current works on factoid TableQA are mainly categorized into two groups: (1) pipeline-based methods consisting of two stages, i.e., cell retrieval and answer reader (Zhu et al., 2021; Chen et al., 2020a); and (2) end-to-end neural networks such as a paradigm of sequence-to-sequence model that takes the context of question answering (e.g., question and table cells) as input to generate natural-language answers (Li et al., 2021b; Pan et al., 2022; Herzig et al., 2021; Pan et al., 2021; Chen, 2023).

Despite much progress made on factoid TableQA, a contradiction between the factoid TableQA and TableQA exists in real scenarios. In factoid TableQA, the answers are always in a short

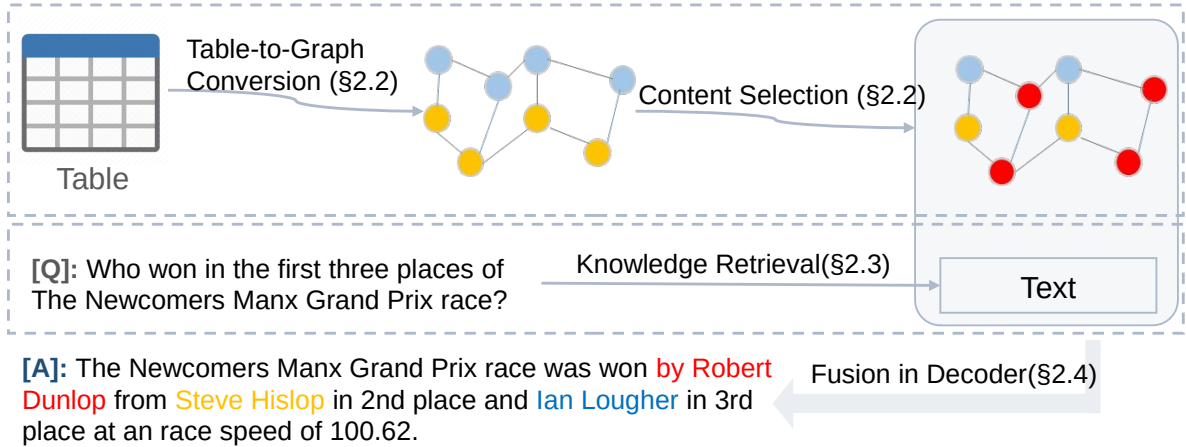


Figure 2: An overview of TAG-QA. The input to TAG-QA is a combination of one table and question, while the output is an answer. The top box shows the content selection process which first converts the table to a graph and selects relevant nodes using GNN. The middle box shows the process of using the sparse retrieval technique to retrieve relevant text as complementary information. The rightmost blue box is to integrate the selected cells and retrieved texts to generate the final answer.

form with a few words directly copied from the relevant table cells. However, in real-world scenarios, the answers are expected to be long and informative sentences in a free form, motivating us to target the free-form TableQA in this paper.

It is challenging to generate coherent and faithful free-form answers over tables. (1) *The well-preserved spatial structure of tables is critical for retrieving relevant table cells to the question.* Different from factoid TableQA, free-form TableQA with sophisticated question shares less semantic similarities to the table content, while depending more on the spatial structure of tables to infer multiple related cells such that the related cells may be located in a relatively connected area, e.g., from either a few selected rows or columns. (2) *The selected table cells, containing the key point, are insufficient for composing the entire coherent sentences.* To generate fluent natural-language sentences as answers, external information such as the relevant background knowledge about the question is necessary. (3) *It is expected to aggregate and reason from the question, retrieved table cells, and external knowledge to compose a reasonable answer.* Given the heterogeneous information, a practical model should be capable of aggregating the information efficiently and generating a coherent and fluent free-form answer.

Figure 1 provides a motivating example to illustrate the insights of this paper. Given a table describing “the 1983 Manx Grand Prix Newcomers Junior Race Results” and a question “Who won

in the first three places of The Newcomers Manx Grand Prix race?”, the goal is to select relevant cells first and then generate a natural sentence as an answer. From this table, we can observe that the state-of-the-art model TAPAS and MATE only select the “rider” while missing the “rank” column, providing low cell selection coverage. For the overall generation quality, we can observe that both the end-to-end T5 (Raffel et al., 2020) and the pipeline-based TAPAS (Herzig et al., 2020) and MATE (Eisenschlos et al., 2021) are missing key information from the table by merely mentioning part of the three riders. In addition, the TAPAS introduces a hallucinated rider named “Spaniard”. These observations motivate us to design a model that can select the relevant cells more accurately and generate faithful answers grounded on the table given a question.

Based on the aforementioned insights, this paper designs a three-stage pipeline framework to tackle the problem of free-form TableQA. Even though the end-to-end TableQA models with high accuracy are prevalently ascribed to the suppression of error accumulated from one-stage training, the long table distracts the model from focusing on relevant table cells, resulting in irrelevant answers. On the other hand, the cell selection module provides a controllable and explainable perspective by extracting a small number of table cells as anchors for the model to generate answers. For the content selection stage, inspired by the recent success of graph models, we convert the table to a graph by

designing the node linking and applying a Graph Neural Network (GNN) to aggregate node information and classify whether the table cell is relevant or not. In addition, to generate informative free-form answers, we employ a sparse retrieval technique to explore extra knowledge from Wikipedia. Consequently, both the extra knowledge and relevant cells are taken into account to calibrate the pre-trained language model bias. Lastly, we adopt a fusion layer in the decoder to generate the final answer.

To summarize, the primary contributions of this paper are three-fold. (1) To the best of our knowledge, we are the first to convert a semi-structured table into a graph, and then design a graph neural network to retrieve relevant table cells. (2) External knowledge is leveraged to fill in the gap between the selected table cell and the long informative answer by providing background information. (3) Comprehensive experiments on a public dataset named FeTaQA (Nan et al., 2022) are performed to verify the effectiveness of TAG-QA. Experimental results show that TAG-QA outperforms the strong baseline TAPAS by 17% and 14%, and outperforms the end-to-end T5 model by 16% and 12%, in terms of BLEU-4 and PARENT F-score, respectively.

2 TAG-QA Approach

In this section, we first formulate the problem of TableQA, and introduce the details of our proposed approach TAG-QA.

2.1 Problem Formulation

A free-form question-answering task is formulated as generating an answer a to a question q based on a semi-structured table T including table cell content and table meta information such as column, and row header. Different from the factoid table question answering task with a short answer, the free-form QA aims at generating informative and long answers.

2.2 Overview

Figure 2 illustrates the overall architecture of our proposed TAG-QA, which is composed of three stages, i.e., relevant table cell localization, relevant external knowledge retrieval, and table-text fusion. (1) *Relevant table cell localization*. We first propose a table-to-graph converter to transform a table into a graph which can preserve the table’s spatial information. We think that the graph-based table representation can better assist in selecting relevant

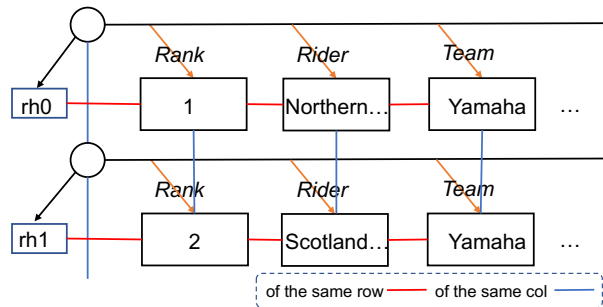


Figure 3: Convert the table shown in Figure 1 into a graphical representation. “ $rh0$ ” is the added row header for the first row. Two relations are considered in the table graph, i.e., “of the same row” and “of the same column” relations.

table cells. (2) *External knowledge retrieval*. We adopt the sparse retrieval technique to collect external information which can be complementary information for the final answer generation. (3) *Table-text fusion*. We employ the fusion-in-decoder model by taking both the selected table cells and the external sources into account to generate the answer. The above three steps enable our model to generate a faithful free-form answer for a question grounded on the table.

2.3 Relevant Table Cell Localization

The initial phase of TAG-QA involves table content selection, a pivotal step that serves as the foundation for subsequent stages. Notably, this stage is of utmost importance as it supplies essential input to the subsequent processes. FeTaQA presents a formidable challenge as a dataset, with a Median/Avg percentage of relevant table cells at 10.7%/16.2%. In order to enhance the precision of the content selection stage, we design a table-to-graph converter to preserve the inherent spatial structure of the tables. We employ GNN to effectively aggregate information at the cell level and subsequently perform a classification task on the table cells.

Table-to-Graph Converter State-of-the-art models prefer to adopt the pre-trained Language Models (LMs) to make predictions by transforming the semi-structured table into natural sentences using a pre-defined template. However, they lose the table structure information and deteriorate the performance of downstream tasks.

TAG-QA designs a table-to-graph converter to transform a table into a graph, preserving the table structure by identifying the cell-to-cell relations.

Figure 3 shows an example of transforming a table into a graph. For the i -th row, we add an empty row header as rhi which reflects the entire row information. All the table cells from the same row are fully connected, and all the table cells from the same column are also fully connected. Besides, we design two types of relations for the table graph, i.e., “of the same row” and “of the same column” relations. In particular, “of the same row” relation captures the entity information, while “of the same column” relation reveals the connection of the same attribute.

In addition, to incorporate the question node into the graph, we create a question node and assign a linking edge between the question and each table cell with the relation “question to cell”.

TAG-QA Content Selection Inspired by QA-GNN (Yasunaga et al., 2021), we propose a content selection module (TAG-CS) that retrieves relevant table cells from the table-based graph. TAG-CS takes the converted table graph from Sec. 2.3 as input, and outputs the question-related table cells. TAG-CS reasons over the table cell level, and each graph node represents a table cell. To fully explore the table semantic and the spatial information, TAG-CS acquires the initial graph node embedding through a pre-trained LM e.g., BERT. Besides, the pre-trained LM and GNN are jointly trained to predict the selected cells.

GNN Architecture We use Graph Attention Network (GAT) (Veličković et al., 2017) which leverages masked self-attention layers and employs iterative message passing among neighbors is applied to predict the selected graph node. GAT follows Eq. 1 to update the i -th node feature $h_i^l \in \mathbb{R}^D$ at layer l through gathering the weighted attention among its neighbors \mathcal{N}_i .

$$h_i^l = f_g \left(\sum_{s \in \mathcal{N}_i \cup \{t\}} \alpha_{st} m_{st} \right) + h_i^{l-1} \quad (1)$$

where α_{st} and $m_{st} \in \mathbb{R}^N$ are the self-attention weight and the message passed from source node s to target node t respectively, and f_g is a 2-layer Multi-Layer Perceptron (MLP) with batch normalization. The message $m_{st} \in \mathbb{R}^N$ from node v_s to v_t is computed using Eq. 2.

$$m_{st} = f_m(h_s^{l-1}, u_s, r_{st}) \quad (2)$$

where $u_s \in \mathbb{R}^{T/2}$ is the source node s feature linearly transformed from the one hot vector node

type u_t . $r_{st} \in \mathbb{R}^T$ is the relation feature from source node s to target node t computed through a 2-layer MLP by taking relation type, source, and target node type into account. f_m is a linear transformation.

The self-attention coefficient α_{st} is updated in Eq. 3. Query and key vectors are linearly transformed by g_q and g_k , as node, edge feature, and the previous layer hidden state provided.

$$\alpha_{st} = \frac{\exp(\gamma_{st})}{\sum_{t' \in \mathcal{N}_s \cup \{s\}} \exp(\gamma_{st'})}, \gamma_{st} = \frac{Q_s^T K_t}{\sqrt{N}} \quad (3)$$

$$Q_s = g_q(h_s^{l-1}, u_s, r_{st}) \quad (4)$$

$$K_t = g_k(h_s^{l-1}, u_t, r_{st}) \quad (5)$$

GNN Training and Inference Given a question q and a table T , TAG-CS reasons over a graph containing both the table cell nodes and the question node by making predictions on the row and column level. We observe that relevant table cells tend to show up in a relatively connected area, thus we make predictions over row and column headers and choose the intersection area. Compared to predicting over the cell level which results in low recall, our method gains a higher chance to capture relevant table cells. For the training stage, TAG-CS maximizes the cross entropy to predict the row and column for relevant cells.

2.4 External Knowledge Retrieval

TAG-QA is the first attempt to leverage the external knowledge to address the table-based free-form QA task. TAG-QA adopts an effective and simple Sparse Retrieval based on the TF/IDF approach to select a potentially relevant context from Wikipedia.

Sparse Retrieval For TAG-QA, the external knowledge is served as a complimentary background context for the next table and text fusion stage. We choose the sparse retrieval method using BM25 (Robertson and Zaragoza, 2009) as a ranking function to retrieve the most relevant text as supplementary information. Given a query q with m keywords k_1, k_2, \dots, k_m , the BM25 ranking score p_i for document d_i is calculated by Eq. 6,

$$p_i = \sum_{j=1}^m \frac{idf(q_j) \times tf(q_j, d_i) \times (\alpha + 1)}{tf(k_j, d_i) + \alpha(1 - \beta + \beta \frac{|d_i|}{L_D})} \quad (6)$$

	Precision	Recall	F-1
TAPAS (Herzig et al., 2020)	65.31	24.20	35.32
MATE (Eisenschlos et al., 2021)	56.93	22.21	31.95
TAG-QA (Ours)	47.60	43.06	45.22

Table 1: Content selection results on FeTaQA dataset.

where idf is the Inverse Document Frequency (IDF), $tf(k_j, d_i)$ is the term frequency of the keyword k_j in document d_i , and L_D is the average document length.

2.5 Table-Text Fusion

After obtaining the predicted highlighted table cells from the table as well as the support context from Wikipedia, TAG-QA aggregates and combines the two information sources through a sequence-to-sequence model Fusion-in-Decoder (FiD) (Izacard and Grave, 2021). FiD appends the question to each information source, encoding each component independently. It subsequently merges all source features and transmits them to the decoder.

Fusion in Decoder Fusion-in-Decoder based on T5 (Raffel et al., 2020) architecture takes question, support context, and the retrieved semi-structured table cells as input. We flatten the highlighted cells as a natural sentence to fit with its pre-trained LM architecture. For the table example shown in Figure 1, the ground-truth selected cells from the first two columns “Rank” and “Rider” can be linearized as “Rank is 1 [SEP] Rider is Northern Ireland Robert D [SEP] Rank is 2 [SEP] Rider is Scotland Steve Hislop [SEP] Rank is 3 [SEP] Rider is Wales Ian Loug.”, where [SEP] is a special token to indicate the end of table slot value.

3 Experiments and Analysis

In this section, we explore the following experimental questions: (1) Does proposed TAG-QA generate a more coherent and faithful answer compared with the baseline? (2) Is table cell selection, knowledge retrieval, and fusion necessary for the free-form TableQA? (3) Is it promising to keep enhancing the three modules of TAG-QA?

3.1 Dataset

This paper focuses on tackling the challenge of generating long free-form answers, rather than the short factoid responses. Consequently, we have opted for the utilization of the state-of-the-art dataset, *FetaQA* (Nan et al., 2022), as our testbed.

The training dataset comprises 7,327 instances, while the development and test sets encompass 1,002 and 2,004 examples, respectively.

3.2 Implementation Details

TAG-CS) TAG-CS applies BERT checkpoint “bert-based-uncased” to learn the table cell representation. For the BERT model, we set the learning rate to $1e-6$ and impose a maximum token length of 35 for each cell. Subsequently, the acquired table cell-level embeddings serve as input node features for our GNN. Within the TAG-CS framework, our GNN module comprises 3 layers, each with node features of 200 dimensions. Additionally, we apply a dropout rate of 0.2 to each layer for regularization.

We train our model on the FeTaQA dataset, configuring it to run for a maximum of 50 epochs. We employ the RAdam optimizer (Liu et al., 2019) with a weight decay of 0.01, utilizing a powerful 24G memory Titan-RTX GPU. To optimize GPU memory usage, we set the maximum number of table cells as 200 and set the batch size as 1. The selection of the best checkpoint is based on the performance of the model on the development set, which is then used for decoding the test set. Additionally, to enhance efficiency, TAG-CS is employed to select intersection cells from the top 3 rows and 3 columns as the relevant cells, drawing upon our accumulated experience in this context. **Sparse Retrieval)** Our implementation relies on the PyTorch-based toolkit Pyserini, designed for reproducible information retrieval research using both sparse and dense representations. We utilize the question as the query to retrieve pertinent contextual information from Wikipedia, selecting the first sentence from the top results. We specifically employ the Lucene Indexes, denoted as “enwiki-paragraphs”².

FiD) In the context of FiD, TAG-QA employs the Adam optimizer with a learning rate of $1e-5$. We select the best checkpoint for inference purposes. In the inference phase, we utilize beam search with a beam size of 3 and apply a length penalty of 1 when generating answers.

3.3 Baselines

To validate the effectiveness of TAG-QA, we choose two different types of methods as baselines, including end-to-end and pipeline-based models.

²<https://github.com/castorini/pyserini>

		BLEU-4	METEOR	ROUGE-L	PARENT (P/R/F)	PARENT-T (P/R/F)
<i>End-to-end</i>						
UniLM	Q-fullTab	17.57	28.30	39.46	38.21/24.18/25.56	26.48/53.99/33.70
	Q-Retrieve	18.46	27.21	39.36	34.12/23.42/23.76	20.37/43.41/25.69
	Q-Retrieve-fullTab	18.89	26.86	38.86	35.29/23.17/23.72	22.07/44.83/27.44
BART	Q-fullTab	7.62	25.70	25.76	39.64/19.68/22.62	25.77/39.53/28.78
	Q-Retrieve	12.20	25.15	28.27	35.55/20.67/22.37	18.31/31.07/20.94
	Q-Retrieve-fullTab	11.97	26.41	28.24	38.45/22.12/23.96	20.57/34.36/23.46
T5	Q-fullTab*	15.66	21.80	35.48	38.88/14.83/18.01	25.11/33.62/26.17
	Q-Retrieve	25.17	24.87	39.89	33.54/20.3/21.68	17.35/31.21/20.13
	Q-Retrieve-fullTab	<u>27.60</u>	<u>26.71</u>	<u>42.38</u>	<u>38.49/23.2/25.06</u>	<u>20.98/35.79/24.02</u>
Oracle-T5	Q-OracleCell	21.77	28.35	42.54	53.37/26.39/30.61	38.21/54.22/41.49
	Q-Retrieve	25.17	24.87	39.89	33.54/20.3/21.68	17.35/31.21/20.13
	Q-Retrieve-OracleCell	31.00	30.35	46.72	46.3/28.44/30.93	27.07/44.32/30.71
<i>Pipeline</i>						
TAPAS-T5	Q-predCell	14.50	21.18	35.51	39.14/12.34/15.67	25.19/29.47/24.38
	Q-Retrieve-predCell	26.81	26.92	42.59	39.23/21.96/24.15	21.43/34.54/23.61
MATE-T5	Q-predCell	14.28	21.01	35.36	39.07/12.2/15.53	24.83/29.56/24.25
	Q-Retrieve-predCell	26.85	26.96	42.60	39.05/21.89/23.99	21.1/34.62/23.57
TAGQA-T5	Q-predCell	17.08	23.22	38.38	41.84/16.53/20.1	27.11/37.03/28.45
	Q-Retrieve-predCell	<u>28.01(↑ 0.41)</u>	<u>27.91(↑ 1.20)</u>	<u>44.16(↑ 1.78)</u>	<u>41.35/23.87/26.2(↑ 1.14)</u>	<u>22.89/37.29/25.64(↑ 1.64)</u>
TAGQA-FiD	Q-Retrieve-predCell*	31.84(↑ 16.18)	30.16(↑ 8.36)	49.39(↑ 13.91)	47.56/26.20/29.59(↑ 11.58)	25.44/39.11/28.26(↑ 2.09)

Table 2: Results on FeTaQA dataset. “P/R/F” denotes the precision/recall/F score. We report end-to-end model UniLM, BART and T5, and the pipeline results. The results of various table cell selection strategies TAPAS, MATE and our proposed TAG with T5 as backbone generation model are noted as TAPAS-T5, MATE-T5 and TagQA-T5. To validate the effectiveness of proposed framework components, we test different combinations of source information to models where “Q” is question, “Retrieve” is the retrieved external knowledge, “fullTab” is full table, and “predCell” refers to the selected table cell. And the last row TAGQA-FiD is the proposed method.

Method	Overall
Reference	4.94
UniLM [end-to-end] (Dong et al., 2019)	3.88
BART [end-to-end] (Lewis et al., 2020)	3.67
T5 [end-to-end] (Raffel et al., 2020)	3.81
Tapas [pipeline] (Herzig et al., 2020)	3.38
MATE [pipeline] (Eisenschlos et al., 2021)	3.30
TAG-QA [pipeline]	3.93

Table 3: Results of human evaluation for reference, end-to-end model and pipeline methods. TAG-QA outperforms the pipeline models by a large margin, and achieves performance on par with the strong end-to-end baseline model T5.

Firstly, we compare TAG-QA with strong state-of-the-art end-to-end pre-trained generative LMs. UniLM (Dong et al., 2019), BART (Lewis et al., 2020), and T5 (Radford et al., 2019). For the input format to the end-to-end model, we flatten the table by concatenating special token [SEP] in between different table cells, and concatenate with the question as a natural sentence, e.g. “*question [SEP] flattened table*”. Furthermore, we compare the performance of our proposed model with pipeline-based methods which include two stages: content selection and answer generation. Content selection makes predictions of relevant cells. We choose two table-based pre-training models: TAPAS (Herzig

et al., 2020) and MATE (Eisenschlos et al., 2021). Moreover, T5 is chosen as the baseline model’s answer generation backbone due to the integration capacity for the table cell and retrieved knowledge.

3.4 Automatic Evaluation Metrics

We use various automatic metrics to evaluate the model performance. Due to the pipeline style of TAG-QA, we report two sets of metrics for content selection and answer generation stages respectively. Firstly, to evaluate the retrieval competency of the table semantic parser, we report Precision, Recall, and F1 scores. Besides, to evaluate the answer generation quality, we choose several automatic evaluation metrics, i.e., BLEU-4 (Papineni et al., 2002), ROUGE-L (Lin, 2004) and METEOR (Banerjee and Lavie, 2005), to evaluate the n-gram match between the generated sentence and the reference answer. Considering the limitation that those metric fails to reflect the faithfulness answer to the fact from the table, we report PARENT (Dhingra et al., 2019) and PARENT-T (Wang et al., 2020) score. PARENT score takes the answer matching with both the reference answer and the table information into account, while PARENT-T focuses on the overlap between the generated answer with the corresponding table.

3.5 Results

We first evaluate the TAG-CS content selection stage table semantic parsing results, as shown in Table 1. For the F-1 score, TAG-QA outperforms the strong baseline model TAPAS and MATE by 9.9% and 13.27%. For recall, TAG-QA achieves the best result, demonstrating that TAG-QA retrieves more relevant table cells. For precision, the baseline model outperforms TAG-QA by retrieving fewer cells which includes more relevant cells. However, the low precision and high recall are a trade-off since the relevant cells make a stronger impact on the overall answer generation quality. Thus, we can tolerate a small amount of irrelevant cells and keep the correct cells as many as possible.

In addition, Table 2 shows the measurements of generated answer quality using TAG-QA compared to previous both end-to-end and pipeline-based state-of-the-art models. From overlapping-based metrics BLEU-4, METEOR, and ROUGE-L, TAG-QA outperforms all the end-to-end and pipeline-based models. Specifically, TAG-QA gains 14.27%/1.86%/9.93% more than the best end-to-end model UniLM in “Q-fullTab” while gains 14.76%/8.98%/13.88% in “Q-predCell” setting, more than the best pipeline-based model TAPAS. For faithfulness metric PARENT and PARENT-T, TAG-QA provides the best performance among the pipeline models by outperforming TAPAS on the “Q-predCell” setting by 13.92% and 3.88% on PARENT and PARENT-T. Compared with end-to-end models, TAG-QA gives the best PARENT score while UniLM shows the best result regarding PARENT-T. It’s explainable because TAG-QA incorporates information outside of the table to generate answers, achieving a trade-off between being grounded on the table and synthesizing informative answers.

Furthermore, to answer Question 2 “Are three stages of the framework necessary to generate high-quality answer?”, we conduct an experiment in Table 2 by comparing the T5 model “Q-fullTab” with pipeline methods backend by T5 using “Q-predCell”. The result shows proposed TAG for content selection TAGQA-T5 selecting 7% of table cell outperforms T5 with fullTab. This indicates the table cell selection is necessary since relevant cells provide an anchor to generate high answer generation. Moreover, to investigate the effect of retrieval knowledge, we show results in Table 2 by concatenating “Retrieval” to the input. The retrieval

Model	BLEU	METEOR	PARENT	PARENT-T
TAG-QA	31.84	30.16	29.59	28.26
TAG-QA w/o JT	31.35	29.65	28.93	27.48
TAG-QA w/o SR	18.93	24.95	21.57	27.95
TAG-QA w/o FiD	21.51	24.03	22.46	25.40

Table 4: Ablation study of the proposed model. We examine the ablated mode by removing the Joint Training (JT) of TAG-CS, Sparse Retrieval (SR), and FiD.

knowledge enhances model performance by providing background knowledge. The proposed model TAGQA-T5 provides the best result by integrating retrieval and informative selected cells. Lastly, our fusion module further enhanced the overall performance by aggregating tables and text efficiently.

Last but not least, to answer the question “Is there space to further enhance performance using this framework?”, we conduct an oracle experiment shown in “Oracle-T5”. With the simple Retrieval technique, T5 backend generation, and oracle table cell, the BLEU-4 result is 31%, and PARENT, PARENT-T are over 30%. If a better retrieval and fusion model is used, the model performance can be further boosted.

3.6 Analysis

To further evaluate the quality of generated answer by various state-of-the-art models when compared to the ground-truth answer, we perform an additional human evaluation. Besides, we conduct an ablation study for TAG-QA to validate the three building blocks: jointly training of LM and GNN for TAG-CS, external context retrieved from Wikipedia, and FiD model. Furthermore, a case study is presented which shows different answer qualities produced by various models.

Human Evaluation Following (Nan et al., 2022), we recruit three human annotators who pass the College English Test (CET-6)³ to judge the quality of the generated sentence. We randomly draw 100 samples from test examples in FeTaQA dataset and collect answers from TAG-QA and baseline models. Then, we present the generated answers to three human annotators without revealing the name of the model, thus reducing human variance.

We provide instructions for human raters to evaluate the sentence quality from four aspects: faithfulness, fluency, correctness, and adequateness. For each aspect, an annotator is supposed to assign a score ranging from 1 (worst) to 5 (best) based on the answer quality. The “overall” column refers to

³A national English as a foreign language test in China.

the average ranking of the model. First, for fluency, the annotator checks if an answer is natural and grammatical. Second, for correctness, we compare the answer with the ground truth by checking if the predicted answer contains the correct information. Third, adequacy reflects if an answer contains all the aspects that are asked. Finally, faithfulness evaluates faithfulness if an answer is faithful and grounded to the contents of the highlighted table region such that it covers all the relevant information from the table while not including other key information outside of the table. From Table 3, we can see TAG-QA ranked the top among all models.

Ablation Study To figure out which building blocks are driving the improvements, we examine different ablated models to understand each component of TAG-QA, including joint training of BERT and GNN from TAG-CS, sparse retrieval, and FiD. Table 4 presents the ablation results under different evaluation metrics. We can see that the model performance drops when any component is removed. Especially, ablating the sparse retrieval module results in the most drop in BLEU-4 and PARENT scores, while removing FiD causes the most significant drop in PARENT-T.

Case Study To inspect the effect of TAG-QA directly, we present a case study in Figure 4, where a sampled table, question, ground-truth relevant table cells (highlighted in blue), the predicted answers of models, as well as the reference are provided. First, we find that the end-to-end model generally contains more information than pipeline models due to the more abundant table information while they suffer from hallucination. For example, T5 and BART identify the ranking position of “*Leandro de Oliveira*” as “17th” while it should be “73rd” from the table. Second, for pipeline models, they tend to generate irrelevant information e.g. MATE mentions the duration and points instead of answering the ranking position and the event. Third, both the end-to-end and pipeline models (TAPAS) fail to cover all the relevant information from the table, e.g. UniLM did not capture the event 12km, and TAPAS fails to mention the position 73rd. By contrast, TAG-QA provides the highest table coverage while keeping the fluency of sentences.

4 Related Work

In this section, we review the related work to ours from the perspectives of TableQA, GNN for natural

Year	Competition	Venue	Pos	Event
RB	Representing Brazil	Representing Brazil	RB	RB
2011	World Cross Country Championships	Punta Umbria Spain	73rd	12 km
2011	World Cross Country Championships	Punta Umbria Spain	17th	Team - 12 km
2011	South American Road Mile Championships	Belém Brazil	1st	One mile
2011	South American Championships	Buenos Aires Argentina	1st	1500m
2011	Pan American Games	Guadalajara México	18th	200m

[Q]: What country did Leandro de Oliveira represent at the 2011 World Cross Country Championships and how did he place?
 [T5-end-to-end]: Leandro de Oliveira represented Brazil at the 2011 World Cross Country Championships and placed 17th.
 [Bart-end-to-end]: Landro de oliveira represented brazil at the 2011 World cross country championships in the 12 km and finished 17th.
 [Unilm-end-to-end]: Leandro de Oliveira was representing Brazil at the 2011 World Cross Country Championships in Punta Umbria , Spain , finishing 73rd .
 [TAPAS-pipeline]: Leandro de Oliveira represented Brazil at the 2011 World Cross Country Championships.
 [MATE-pipeline]: Leandro de Oliveira finished in 37:10 and finished with 326 points.
 [Ours-pipeline]: Leandro de Oliveira represented Brazil at the 2011 World Cross Country Championships and placed 73rd in the 12 km race.
 [REF]: Leandro de Oliveira represented Brazil at the 2011 World Cross Country Championships and placed 73rd in the 12 km race.

Figure 4: A case study from FeTaQA for qualitative analysis. The highlighted cells are the ground-truth relevant table cells. “RB” refers to “Representing Brazil”. Hallucinated content from the predicted answer is marked in red and the correct content in blue.

language processing, and knowledge-grounded text generation.

TableQA FeTaQA is the first TableQA dataset that addresses the significance of free-form answer generation, while most current research work including WikiTableQuestions (Pasupat and Liang, 2015), Spider (Yu et al., 2018), HybridQA (Chen et al., 2020b), OTT-QA (Chen et al., 2020a), and TAT-QA (Zhu et al., 2021) focuses on the short facitoid answer generation. The early solution (Zhong et al., 2017; Liang et al., 2017) of addressing the TableQA is to parse the natural question into a machine-executable meaning representations that can be used to query the table. To reduce the labor-intensive logical annotation, a semantic parser trained over weak supervision from denotations has been drawing attention. Plenty of Transformer-based table pre-training models demonstrate decent TableQA performance, e.g., TaPas (Herzig et al., 2020), MATE (Eisenschlos et al., 2021), TaBERT (Yin et al., 2020), StruG (Deng et al., 2021), GraPPa (Yu et al., 2021), and TaPEX (Liu et al., 2022a). In addition, rather than explore table structure, RCI (Glass et al., 2021) assumes the row and column are independent, and predicts the probability of containing the answer to a question in each row and column of a table individually.

GNN for Natural Language Processing Apart from the extensively renowned causal language

models that have showcased impressive results in various task (Vaswani et al., 2017; Parmar et al., 2018; Wang et al., 2023a, 2022, 2023b), a rich variety of language processing tasks gain improvements from exploiting the power of GNN (Li et al., 2015). Tasks such as semantic parsing (Chen et al., 2021a), text classification (Lin et al., 2021), text generation (Fei et al., 2021), question answering (Wang et al., 2021; Yasunaga et al., 2021) can be expressed with a graph structure and handled with graph-based methods. In addition, researchers apply GNN to model the text generation from structured data tasks e.g. graph-to-sequence (Marchegiani and Perez-Beltrachini, 2018), and AMR-to-text (Ribeiro et al., 2019).

Knowledge-Grounded Text Generation

Encoder-decoder-based models have been proposed to tackle the generation task by mapping the input to the output sequence. However, the input text is insufficient to provide knowledge to generate decent output due to the lack of commonsense, factual events, and semantic information. Knowledge-grounded text generation incorporating external knowledge such as linguistic features (Liu et al., 2021c), knowledge graph (Liu et al., 2021b; Li et al., 2021a), knowledge base (Eric and Manning, 2017; He et al., 2017; Liu et al., 2022b), and textual knowledge (Liu et al., 2021a; Zhao et al., 2021) help to generate a more logical and informative answer.

5 Conclusion

This paper presents a generalized pipeline-based framework TAG-QA for free-form long answer generation for TableQA. The core idea of TAG-QA is to divide the answer generation process into three stages: (1) transform the table into a graph and jointly reason over the question-table graph to select relevant cells; (2) retrieve contextual knowledge from Wikipedia using sparse retrieval, and (3) integrate the selected cells with the content knowledge to predict the final answer. Extensive experiments on a public dataset FeTaQA are conducted to verify the generated answer quality from both the fluency and faithfulness aspects.

Limitations

One limitation of TAG-CS, which accepts the entire table as input, arises when dealing with large tables, as training both BERT and the graph model

simultaneously becomes challenging due to GPU memory constraints. Consequently, one promising avenue for future research involves the efficient modeling of large tables. Furthermore, it’s worth noting that the availability of only one public dataset, FeTaQA, for free-form TableQA, has constrained our validation efforts to this single dataset. However, we are committed to expanding the scope of our research in the future by evaluating the performance of our pipeline model, TAG-QA, across multiple free-form TableQA datasets.

References

- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Robin D Burke, Kristian J Hammond, Vladimir Kulyukin, Steven L Lytinen, Noriko Tomuro, and Scott Schoenberg. 1997. Question answering from frequently asked question files: Experiences with the faq finder system. *AI magazine*, 18(2):57–57.
- Wenhu Chen. 2023. [Large language models are few\(1\)-shot table reasoners](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1120–1130, Dubrovnik, Croatia. Association for Computational Linguistics.
- Wenhu Chen, Ming-Wei Chang, Eva Schlinger, William Wang, and William W Cohen. 2020a. Open question answering over tables and text. *arXiv preprint arXiv:2010.10439*.
- Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Wang. 2020b. Hybridqa: A dataset of multi-hop question answering over tabular and textual data. *arXiv preprint arXiv:2004.07347*.
- Zhi Chen, Lu Chen, Yanbin Zhao, Ruisheng Cao, Zihan Xu, Su Zhu, and Kai Yu. 2021a. [ShadowGNN: Graph projection neural network for text-to-SQL parser](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5567–5577, Online. Association for Computational Linguistics.
- Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. 2021b. [FinQA: A dataset of numerical reasoning over financial data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3697–3711, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Xiang Deng, Ahmed Hassan Awadallah, Christopher Meek, Oleksandr Polozov, Huan Sun, and Matthew Richardson. 2021. [Structure-grounded pretraining for text-to-SQL](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1337–1350, Online. Association for Computational Linguistics.
- Bhuwan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William Cohen. 2019. [Handling divergent reference texts when evaluating table-to-text generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4884–4895, Florence, Italy. Association for Computational Linguistics.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. [Unified language model pre-training for natural language understanding and generation](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Julian Eisenschlos, Maharshi Gor, Thomas Müller, and William Cohen. 2021. [MATE: Multi-view attention for table transformer efficiency](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7606–7619, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mihail Eric and Christopher Manning. 2017. [A copy-augmented sequence-to-sequence architecture gives good performance on task-oriented dialogue](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 468–473, Valencia, Spain. Association for Computational Linguistics.
- Zichu Fei, Qi Zhang, and Yaqian Zhou. 2021. [Iterative GNN-based decoder for question generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2573–2582, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Michael Glass, Mustafa Caim, Alfio Gliozzo, Saneem Chemmengath, Vishwajeet Kumar, Rishav Chakravarti, Avi Sil, Feifei Pan, Samarth Bharadwaj, 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, Online. Association for Computational Linguistics.
- Yanchao Hao, Yuanzhe Zhang, Kang Liu, Shizhu He, Zhanyi Liu, Hua Wu, and Jun Zhao. 2017. [An end-to-end model for question answering over knowledge base with cross-attention combining global knowledge](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 221–231, Vancouver, Canada. Association for Computational Linguistics.
- Shizhu He, Cao Liu, Kang Liu, and Jun Zhao. 2017. [Generating natural answers by incorporating copying and retrieving mechanisms in sequence-to-sequence learning](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 199–208, Vancouver, Canada. Association for Computational Linguistics.
- Jonathan Herzig, Thomas Müller, Syrine Krichene, and Julian Eisenschlos. 2021. [Open domain question answering over tables via dense retrieval](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 512–519, Online. Association for Computational Linguistics.
- Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Eisenschlos. 2020. [TaPas: Weakly supervised table parsing via pre-training](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4320–4333, Online. Association for Computational Linguistics.
- Gautier Izacard and Edouard Grave. 2021. [Leveraging passage retrieval with generative models for open domain question answering](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880, Online. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Junyi Li, Wayne Xin Zhao, Zhicheng Wei, Nicholas Jing Yuan, and Ji-Rong Wen. 2021a. [Knowledge-based review generation by coherence enhanced text planning](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 183–192.
- Xiao Li, Yawei Sun, and Gong Cheng. 2021b. [Tsqa: tabular scenario based question answering](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13297–13305.
- Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. 2015. [Gated graph sequence neural networks](#). *arXiv preprint arXiv:1511.05493*.
- Chen Liang, Jonathan Berant, Quoc Le, Kenneth D. 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, Vancouver, Canada. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yuxiao Lin, Yuxian Meng, Xiaofei Sun, Qinghong Han, Kun Kuang, Jiwei Li, and Fei Wu. 2021. Bertgcn: Transductive text classification by combining gcn and bert. *arXiv preprint arXiv:2105.05727*.
- Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. 2019. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*.
- Qian Liu, Bei Chen, Jiaqi Guo, Morteza Ziyadi, Zeqi Lin, Weizhu Chen, and Jian-Guang Lou. 2022a. TAPEX: Table pre-training via learning a neural SQL executor. In *International Conference on Learning Representations*.
- Ye Liu, Kazuma Hashimoto, Yingbo Zhou, Semih Yavuz, Caiming Xiong, and Philip Yu. 2021a. Dense hierarchical retrieval for open-domain question answering. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ye Liu, Yao Wan, Lifang He, Hao Peng, and Philip S Yu. 2021b. Kg-bart: Knowledge graph-augmented bart for generative commonsense reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6418–6425.
- Ye Liu, Yao Wan, Jianguo Zhang, Wenting Zhao, and Philip Yu. 2021c. Enriching non-autoregressive transformer with syntactic and semantic structures for neural machine translation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1235–1244, Online. Association for Computational Linguistics.
- Ye Liu, Semih Yavuz, Rui Meng, Dragomir Radev, Caiming Xiong, and Yingbo Zhou. 2022b. Uni-parser: Unified semantic parser for question answering on knowledge base and database. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8858–8869, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Diego Marcheggiani and Laura Perez-Beltrachini. 2018. Deep graph convolutional encoders for structured data to text generation. In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 1–9, Tilburg University, The Netherlands. Association for Computational Linguistics.
- Linyong Nan, Chiachun Hsieh, Ziming Mao, Xi Victoria Lin, Neha Verma, Rui Zhang, Wojciech Kryściński, Hailey Schoelkopf, Riley Kong, Xiangru Tang, Mutethia Mutuma, Ben Rosand, Isabel Trindade, Renusree Bandaru, Jacob Cunningham, Caiming Xiong, Dragomir Radev, and Dragomir Radev. 2022. FeTaQA: Free-form Table Question Answering. *Transactions of the Association for Computational Linguistics*, 10:35–49.
- Feifei Pan, Mustafa Canim, Michael Glass, Alfio Gliozzo, and Peter Fox. 2021. CLTR: An end-to-end, transformer-based system for cell-level table retrieval and table question answering. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 202–209, Online. Association for Computational Linguistics.
- Feifei Pan, Mustafa Canim, Michael Glass, Alfio Gliozzo, and James Hendler. 2022. End-to-end table question answering via retrieval-augmented generation. *arXiv preprint arXiv:2203.16714*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. 2018. Image transformer. In *International conference on machine learning*, pages 4055–4064. PMLR.
- Panupong Pasupat and Percy Liang. 2015. Compositional semantic parsing on semi-structured tables. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1470–1480, Beijing, China. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- 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. *Journal of Machine Learning Research*, 21(140):1–67.
- Leonardo F. R. Ribeiro, Claire Gardent, and Iryna Gurevych. 2019. Enhancing AMR-to-text generation with dual graph representations. 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 3183–3194, Hong Kong, China. Association for Computational Linguistics.

- Stephen Robertson and Hugo Zaragoza. 2009. *The probabilistic relevance framework: BM25 and beyond*. Now Publishers Inc.
- Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang, Akari Asai, Gabriel Ilharco, Hananeh Hajishirzi, and Jonathan Berant. 2021. Multimodalqa: Complex question answering over text, tables and images. *arXiv preprint arXiv:2104.06039*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *International Conference on Learning Representations (ICLR)*.
- Kuan Wang, Yuyu Zhang, Diyi Yang, Le Song, and Tao Qin. 2021. Gnn is a counter? revisiting gnn for question answering. *arXiv preprint arXiv:2110.03192*.
- Yibo Wang, Congying Xia, Guan Wang, and Philip S. Yu. 2022. Continuous prompt tuning based textual entailment model for e-commerce entity typing. In *2022 IEEE International Conference on Big Data (Big Data)*, pages 1383–1388.
- Yibo Wang, Yanbing Xue, Bo Liu, Musen Wen, Wenting Zhao, Stephen Guo, and Philip S Yu. 2023a. Click-conversion multi-task model with position bias mitigation for sponsored search in ecommerce. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1884–1888.
- Yibo Wang, Wenting Zhao, Yao Wan, Zhongfen Deng, and Philip S. Yu. 2023b. Named entity recognition via machine reading comprehension: A multi-task learning approach. *arXiv:2309.11027*.
- Zhenyi Wang, Xiaoyang Wang, Bang An, Dong Yu, and Changyou Chen. 2020. **Towards faithful neural table-to-text generation with content-matching constraints**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1072–1086, Online. Association for Computational Linguistics.
- Xuchen Yao and Benjamin Van Durme. 2014. **Information extraction over structured data: Question answering with Freebase**. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 956–966, Baltimore, Maryland. Association for Computational Linguistics.
- Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, and Jure Leskovec. 2021. **QA-GNN: Reasoning with language models and knowledge graphs for question answering**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 535–546, Online. Association for Computational Linguistics.
- Pengcheng Yin, Graham Neubig, Wen tau Yih, and Sebastian Riedel. 2020. TaBERT: Pretraining for joint understanding of textual and tabular data. In *Annual Conference of the Association for Computational Linguistics (ACL)*.
- Tao Yu, Chien-Sheng Wu, Xi Victoria Lin, bailin wang, Yi Chern Tan, Xinyi Yang, Dragomir Radev, richard socher, and Caiming Xiong. 2021. **Gra{pp}a: Grammar-augmented pre-training for table semantic parsing**. In *International Conference on Learning Representations*.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018. **Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921, Brussels, Belgium. Association for Computational Linguistics.
- Wenting Zhao, Ye Liu, Yao Wan, and Philip Yu. 2021. **Attend, memorize and generate: Towards faithful table-to-text generation in few shots**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4106–4117, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*.
- Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. **TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3277–3287, Online. Association for Computational Linguistics.