

HybridQA: A Dataset of Multi-Hop Question Answering over Tabular and Textual Data

Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, William Wang

University of California, Santa Barbara, CA, USA

{wenhuchen, hwzha, xwhan}@cs.ucsb.edu,

{zhiyuchen, hongwang600, william}@cs.ucsb.edu

Abstract

Existing question answering datasets focus on dealing with homogeneous information, based either only on text or KB/Table information alone. However, as human knowledge is distributed over heterogeneous forms, using homogeneous information alone might lead to severe coverage problems. To fill in the gap, we present HybridQA¹, a new large-scale question-answering dataset that requires reasoning on heterogeneous information. Each question is aligned with a Wikipedia table and multiple free-form corpora linked with the entities in the table. The questions are designed to aggregate both tabular information and text information, i.e., lack of either form would render the question unanswerable. We test with three different models: 1) a table-only model. 2) text-only model. 3) a hybrid model that combines heterogeneous information to find the answer. The experimental results show that the EM scores obtained by two baselines are below 20%, while the hybrid model can achieve an EM over 40%. This gap suggests the necessity to aggregate heterogeneous information in HybridQA. However, the hybrid model’s score is still far behind human performance. Hence, HybridQA can serve as a challenging benchmark to study question answering with heterogeneous information.

1 Introduction

Question answering systems aim to answer any form of question of our interests, with evidence provided by either free-form text like Wikipedia passages (Rajpurkar et al., 2016; Chen et al., 2017; Yang et al., 2018) or structured data like Freebase/WikiData (Berant et al., 2013; Kwiatkowski et al., 2013; Yih et al., 2015; Weston et al., 2015) and WikiTables (Pasupat and Liang, 2015). Both

forms have their advantages, the free-form corpus has in general better coverage while structured data has better compositionality to handle complex multi-hop questions. Due to the advantages of different representation forms, people like to combine them in real world applications. Therefore, it is sometime not ideal to assume the question has answer in a passage. This paper aims to simulate a more realistic setting where the evidences are distributed into heterogeneous data, and the model requires to aggregate information from different forms for answering a question. There has been some pioneering work on building hybrid QA systems (Sun et al., 2019, 2018; Xiong et al., 2019). These methods adopts KB-only datasets (Berant et al., 2013; Yih et al., 2015; Talmor and Berant, 2018) to simulate a hybrid setting by randomly masking KB triples and replace them with text corpus. Experimental results have proved decent improvement, which shed lights on the potential of hybrid question answering systems to integrate heterogeneous information.

Though there already exist numerous valuable questions answering datasets as listed in Table 1, these datasets were initially designed to use either structured or unstructured information during annotation. There is no guarantee that these questions need to aggregate heterogeneous information to find the answer. Therefore, designing hybrid question answering systems would probably yield marginal benefits over the non-hybrid ones, which greatly hinders the research development in building hybrid question answering systems.

To fill in the gap and construct a heterogeneous QA dataset, we present HYBRIDQA, which is collected by crowdsourcing based on Wikipedia tables. During annotation, each crowd worker is presented with a table along with its hyperlinked Wikipedia passages to propose questions requiring multi-hop reasoning over both forms of informa-

¹<https://github.com/wenhuchen/HybridQA>

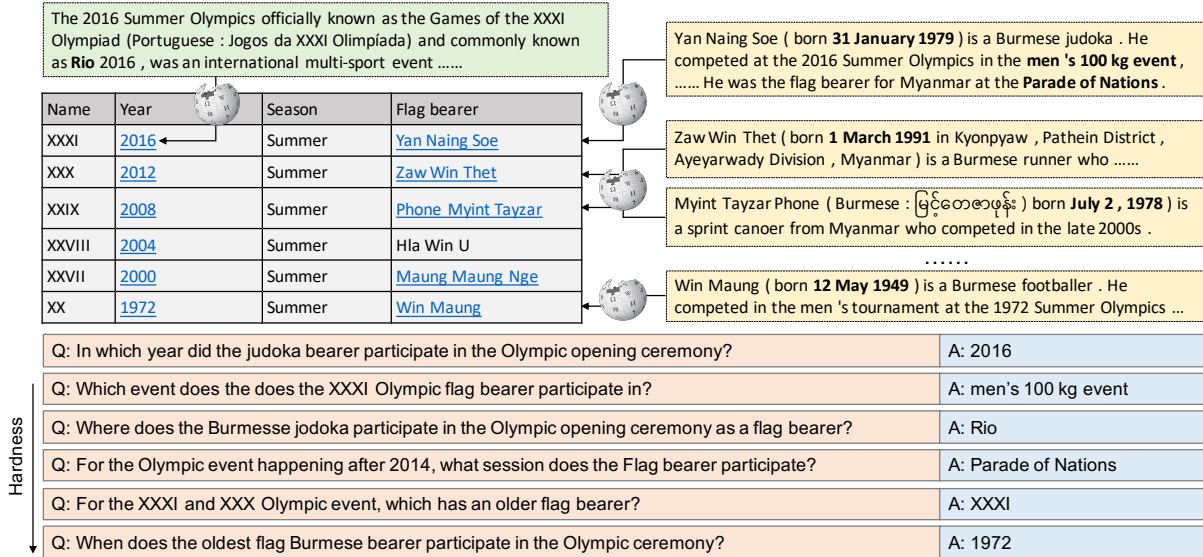


Figure 1: Examples of annotated question answering pairs from Wikipedia page². Underlined entities have hyper-linked passages, which are displayed in the boxes. The lower part shows the human-annotated question-answer pairs roughly categorized based on their hardness.

Dataset	Size	#Docs	KB/ Table	Multi-Hop
WebQuestions	5.8K	no	yes	yes
WebQSP	4.7K	no	yes	yes
WebQComplex	34K	no	yes	yes
MetaQA	400k	no	yes	yes
WikiTableQA	22K	no	yes	yes
SQuAD-v1	107K	1	no	no
DROP	99K	1	no	yes
TriviaQA	95K	>1	no	no
HotpotQA	112K	>1	no	yes
Natural-QA	300K	>1	no	yes
HYBRIDQA	70K	>>1	yes	yes

Table 1: Comparison of existing datasets, where #docs means the number of documents provided for a specific question. 1) KB-only datasets: WebQuestions (Berant et al., 2013), WebQSP (Yih et al., 2016), WebComplex (Talmor and Berant, 2018), MetaQA (Zhang et al., 2018), WikiTableQuestion (Pasupat and Liang, 2015). 2) Text-only datasets with single passage: like SQuAD (Rajpurkar et al., 2016), DROP (Dua et al., 2019). 3) open-domain Text-Only dataset: TriviaQA (Joshi et al., 2017), HotpotQA (Yang et al., 2018), Natural Questions (Kwiatkowski et al., 2019).

tion. The dataset consists of roughly 70K question-answering pairs aligned with 13,000 Wikipedia tables. As Wikitable (Bhagavatula et al., 2013) are curated from high-standard professionals to organize a set of information regarding a given theme, its information is mostly absent in the text. Such a complementary nature makes Wikitable an ideal environment for hybrid question answering. To ensure that the answers cannot be hacked by single-

#Table	#Row/#Column	#Cell	#Links/Table
13,000	15.7/4.4	70	44
#Passage	#Words/Passage	#Ques	#Words/Ques
293,269	103	69,611	18.9

Table 2: Statistics of Table and Passage in our dataset.

hop or homogeneous models, we carefully employ different strategies to calibrate the annotation process. An example is demonstrated in Figure 1. This table is aimed to describe Burmese flag bearers over different Olympic events, where the second column has hyperlinked passages about the Olympic event, and the fourth column has hyperlinked passages about biography individual bearers. The dataset is both multi-hop and hybrid in the following senses: 1) the question requires multiple hops to achieve the answer, each reasoning hop may utilize either tabular or textual information. 2) the answer may come from either the table or a passage.

In our experiments, we implement three models, namely Table-only model, Passage-only, and a heterogeneous model HYBRIDER, which combines both information forms to perform multi-hop reasoning. Our Experiments show that two homogeneous models only achieve EM lower than 20%, while HYBRIDER can achieve an EM over 40%, which concludes the necessity to do multi-hop reasoning over heterogeneous information on HYBRIDQA. As the HYBRIDER is still far behind human performance, we believe that it would be a challenging next-problem for the community.

2 Dataset

In this section, we describe how we crawl high-quality tables with their associated passages, and then describe how we collect hybrid questions. The statistics of HYBRIDQA is in Table 2.

Table/Passage Collection To ease the annotation process, we apply the following rules during table crawling. 1) we need tables with rows between 5-20, columns between 3-6, which is appropriate for the crowd-workers to view. 2) we restrain the tables from having hyperlinked cells over 35% of its total cells, which provide an abundant amount of textual information. For each hyperlink in the table, we retrieve its Wikipedia page and crop at most the first 12 sentences from its introduction session as the associated passage. 3) we apply some additional rules to avoid improper tables and finally collect 13,000 high-quality tables.

Question/Answer Collection We release 13K HITs (human intelligence task) on the Amazon Mechanical Turk platform, where each HIT presents the crowd-worker with one crawled Wikipedia table along with its hyperlinked passages. We require the worker to write six questions as well as their answers. The question annotation phase is not trivial, as we specifically need questions that rely on both tabular and textual information. In order to achieve that, we exemplify abundant examples in our Amazon Turker interface with detailed explanations to help crowd-workers to understand the essence of the “hybrid” question. The guidelines are described as follows:

- The question requires multiple steps over two information forms of reasoning to answer.
- Table reasoning step specifically includes (i) filter our table rows based on equal/greater/less, e.g. “For the XXXI Olympic event”, (ii) superlative operation over a column, e.g. “the earliest Olympic event”, (iii) hop between two cells, e.g. “Which event ... participate in ...”, (iv) extract information from table, e.g. “In which year did the player ...”.
- Text reasoning step specifically includes (i) select passages based on the certain mentions, e.g. “the judoka bearer”, (ii) extract a span from the passage as the answer.
- The answer should be a minimum text span from either a table cell or a specific passage.

Based on the above criteria, we hire five CS-majored graduate students as our “human expert” to decide the acceptance of a HIT. The average completion time for one HIT is 12 minutes, and payment is \$2.3 U.S. dollars/HIT.

Annotation De-biasing As has been suggested in previous papers (Kaushik and Lipton, 2018; Chen and Durrett, 2019; Clark et al., 2019), the existing benchmarks on multi-hop reasoning question answering have annotation biases, which makes designing multi-hop models unnecessary. We discuss different biases and our prevention as follows:

- *Table Bias*: our preliminary study observes that the annotators prefer to ask questions regarding the top part of the table. In order to deal with this issue, we explicitly highlight certain regions in the table to encourage crowd-workers to raise questions regarding the given uniformly-distributed regions.
- *Passage Bias*: the preliminary study shows that the annotators like to ask questions regarding the first few sentences in the passage. In order to deal with such a bias, we use an algorithm to match the answer with linked passages to find their span and reject the HITs, which have all the answers centered around the first few sentences.
- *Question Bias*: the most difficult bias to deal with is the “fake” hybrid question like “when is 2012 Olympic Burmese runner flag bearer born?” for the table listed in Figure 1. Though it seems that “2012 Olympic” is needed to perform hop operation on the table, the “runner flag bearer” already reviews the bearer as “Zaw Win Thet” because there is no other runner bearer in the table. With that said, reading the passage of “Zaw Win Thet” alone can simply lead to the answer. In order to cope with such a bias, we ask “human experts” to spot such questions and reject them.

Statistics After we harvest the human annotations from 13K HITs (78K questions), we trace back the answers to its source (table or passage). Then we apply several rules to further filter out low-quality annotations: 1) the answer cannot be found from either table or passage, 2) the answer is longer

than 20 words, 3) using a TF-IDF retriever can directly find the answer passage with high similarity without relying on tabular information.

We filter the question-answer pairs based on the previous criteria and release the filtered version. As our goal is to solve multi-hop hybrid questions requiring a deeper understanding of heterogeneous information. We follow HotpotQA (Yang et al., 2018) to construct a more challenging dev/test split in our benchmark. Specifically, we use some statistical features like the “size of the table”, “similarity between answer passage and question”, “whether question directly mentions the field”, etc. to roughly classify the question into two difficulty levels: simple (65%) and hard (35%). We construct our dev and test set by sampling half-half from the two categories. We match the answer span against all the cells and passages in the table and divide the answer source into three categories: 1) the answer comes from a text span in a table cell, 2) the answer comes from a certain linked passage, 3) the answer is computed by using numerical operation like ‘count’, ‘add’, ‘average’, etc. The matching process is approximated, not guaranteed to be 100% correct. We summarize our findings in Table 3. In the following experiments, we will report the EM/F1 score for these fine-grained question types to better understand our results.

Split	Train	Dev	Test	Total
In-Passage	35,215	2,025	20,45	39,285 (56.4%)
In-Table	26,803	1,349	1,346	29,498 (42.3%)
Computed	664	92	72	828 (1.1%)
Total	62,682	3,466	3,463	69,611

Table 3: Data Split: In-Table means the answer comes from plain text in the table, and In-Passage means the answer comes from certain passage.

3 Data Analysis

In this section, we specifically analyze the different aspects of the dataset to provide the overall characteristics of the new dataset.

3.1 Question Types

We heuristically identified question types for each collected question. To identify the question type, we locate the central question word (CQW) in the question and take the neighboring three tokens (Yang et al., 2018) to determine the question types. We visualize the distribution in Figure 2, which demonstrates the syntactic diversity of the questions in HYBRIDQA.




Figure 2: The type of questions in HYBRIDQA, question types are extracted using rules starting at the question words or preposition before them.

3.2 Answer Types

We further sample 100 examples from the dataset and present the types of answers in Table 4. As can be seen, it covers a wide range of answer types. Compared to (Yang et al., 2018), our dataset covers more number-related or date-related questions, which reflects the nature of tabular data.

Answer Type	%	Example(s)
Location	22	Balestier Road, Atlanta
Number	22	762 metres (2,500 ft), 15,000
Date	20	April 25 , 1994, 1913
Person	15	Bärbel Wöckel, Jerry
Group	3	Hallmark Entertainment
Event	3	Battle of Hlobane
Artwork	1	Songmaster
Adjective	4	second-busiest
Other proper noun	8	Space Opera, CR 131
Common noun	1	other musicians

Table 4: Types of answers in HYBRIDQA.

3.3 Inference Types

analyze multi-hop reasoning types in Figure 3. According to our statistics, most of the questions require two or three hops to find the answer.

1) Type I question (23.4%) uses Table → Passage chain, it first uses table-wise operations (equal/greater/less/first/last/argmax/argmin) to locate certain cells in the table, and then hop to their neighboring hyperlinked cells within the same row, finally extracts a text span from the passage of the

Type I (T->P) Q: Where was the XXXI Olympic held? A: Rio	<table border="1"> <thead> <tr> <th>Name</th> <th>Event year</th> </tr> </thead> <tbody> <tr> <td>XXXI</td> <td>2016</td> </tr> </tbody> </table>	Name	Event year	XXXI	2016	... commonly known as Rio 2016 , was an international multi-sport event					
Name	Event year										
XXXI	2016										
Type II (P->T) Q: What was the name of the Olympic event held in Rio? A: XXXI	<table border="1"> <thead> <tr> <th>Name</th> <th>Event year</th> </tr> </thead> <tbody> <tr> <td>XXXI</td> <td>2016</td> </tr> </tbody> </table>	Name	Event year	XXXI	2016	... commonly known as Rio 2016 , was an international multi-sport event					
Name	Event year										
XXXI	2016										
Type III (P->T->P) Q: When was the flag bearer of Rio Olympic born? A: 31 January 1979	<table border="1"> <thead> <tr> <th>Flag Bearer</th> <th>Event year</th> </tr> </thead> <tbody> <tr> <td>Yan Naing Soe (born 31 January 1979) ...</td> <td>2016</td> </tr> </tbody> </table>	Flag Bearer	Event year	Yan Naing Soe (born 31 January 1979) ...	2016	... commonly known as Rio 2016 , was an international					
Flag Bearer	Event year										
Yan Naing Soe (born 31 January 1979) ...	2016										
Type IV (T&S) Q: Which male bearer participated in Men's 100kg event in the Olympic game? A: Yan Naing Soe	<table border="1"> <tbody> <tr> <td>Yan Naing Soe ... Men's 100kg event</td> <td>Flag Bearer</td> <td>Gender</td> </tr> <tr> <td>Zaw Win Thet ... Men's 400m running</td> <td>Yan Naing Soe</td> <td>Male</td> </tr> <tr> <td></td> <td>Zaw Win Thet</td> <td>Male</td> </tr> </tbody> </table>	Yan Naing Soe ... Men's 100kg event	Flag Bearer	Gender	Zaw Win Thet ... Men's 400m running	Yan Naing Soe	Male		Zaw Win Thet	Male	
Yan Naing Soe ... Men's 100kg event	Flag Bearer	Gender									
Zaw Win Thet ... Men's 400m running	Yan Naing Soe	Male									
	Zaw Win Thet	Male									
Type V (T-Compare P-Compare) Q: For the 2012 and 2016 Olympic Event, when was the younger flag bearer born? A: 1 March 1991	<table border="1"> <thead> <tr> <th>Flag Bearer</th> <th>Event year</th> </tr> </thead> <tbody> <tr> <td>Yan Naing Soe (born 31 January 1979) ...</td> <td>2016</td> </tr> <tr> <td>Zaw Win Thet (born 1 March 1991)</td> <td>2012</td> </tr> </tbody> </table>	Flag Bearer	Event year	Yan Naing Soe (born 31 January 1979) ...	2016	Zaw Win Thet (born 1 March 1991)	2012				
Flag Bearer	Event year										
Yan Naing Soe (born 31 January 1979) ...	2016										
Zaw Win Thet (born 1 March 1991)	2012										
Type VI (T-Superlative P-Superlative) Q: When did the youngest Burmese flag bearer participate in the Olympic opening ceremony? A: 2012	<table border="1"> <thead> <tr> <th>Flag Bearer</th> <th>Event year</th> </tr> </thead> <tbody> <tr> <td>Yan ... 31 January 1979) ...</td> <td>2016</td> </tr> <tr> <td>Zaw ... 1 March 1991) ...</td> <td>2012</td> </tr> <tr> <td>Myint ... July 2, 1978) ...</td> <td>2008</td> </tr> </tbody> </table>	Flag Bearer	Event year	Yan ... 31 January 1979) ...	2016	Zaw ... 1 March 1991) ...	2012	Myint ... July 2, 1978) ...	2008		
Flag Bearer	Event year										
Yan ... 31 January 1979) ...	2016										
Zaw ... 1 March 1991) ...	2012										
Myint ... July 2, 1978) ...	2008										

Figure 3: Illustration of different types of multi-hop questions.

hyperlinked cell as the answer.

2) Type II question (20.3%) uses Passage → Table chain, it first uses cues present in the question to retrieve related passage, which traces back to certain hyperlinked cells in the table, and then hop to a neighboring cell within the same row, finally extracts text span from that cell.

3) Type III question (35.1%) uses Passage → Table → Passage chain, it follows the same pattern as Type II, but in the last step, it hops to a hyperlinked cell and extracts answer from its linked passage. This is the most common pattern.

4) Type IV question (17.3%) uses Passage and Table jointly to identify a hyperlinked cell based on table operations and passage similarity and then extract the plain text from that cell as the answer.

5) Type V question (3.1%) involves two parallel reasoning chain, while the comparison is involved in the intermediate step to find the answer.

6) Type VI questions (0.8%) involve multiple reasoning chains, while superlative is involved in the intermediate step to obtain the correct answer.

4 Model

In this section, we propose three models we use to perform question answering on HYBRIDQA.

4.1 Table-Only Model

In this setting, we design a model that can only rely on the tabular information to find the answer. Our model is based on the SQL semantic parser (Zhong

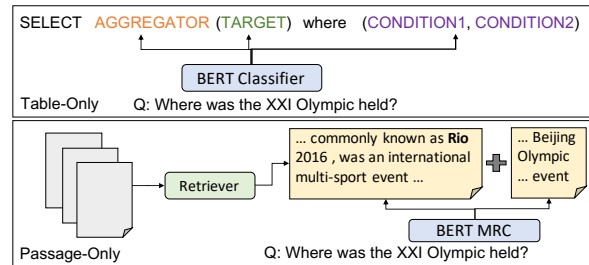


Figure 4: Illustration of the table-only and passage-only baselines, both are based on BERT Encoder.

et al., 2017; Xu et al., 2017), which uses a neural network to parse the given questions into a symbolic form and execute against the table. We follow the SQLNet (Xu et al., 2017) to flatten the prediction of the whole SQL query into a slot filling procedure. More specifically, our parser model first encode the input question q using BERT (Devlin et al., 2019) and then decode the aggregation, target, condition separately as described in Figure 4. The aggregation slot can have the following values of “argmax, argmin, argmax-date, argmin-date”, the target and condition slots have their potential values based on the table field and its corresponding entries. Though we do not have the ground-truth annotation for these simple SQL queries, we can use heuristics to infer them from the denotation. We use the synthesized question-SQL pairs to train the parser model.

4.2 Passage-Only Model

In this setting, we design a model that only uses the hyperlinked passages from the given table to find the answer. Our model is based on DrQA (Chen et al., 2017), which first uses an ensemble of several retrievers to retrieve related documents and then concatenate several documents together to do reading comprehension with the state-of-the-art BERT model (Devlin et al., 2019). The basic architecture is depicted in Figure 4, where we use the retriever to retrieve the top-5 passages from the pool and then concatenate them as a document for the MRC model, and the maximum length of the concatenated document is set to 512.

4.3 HYBRIDER

In order to cope with heterogeneous information, we propose a novel architecture called HYBRIDER. We divide the model into two phases as depicted in Figure 6 and describe them separately below:

Linking This phase is aimed to link questions to their related cells from two sources:

- Cell Matching: it aims to link cells explicitly mentioned by the question. The linking consists of three criteria, 1) the cell’s value is explicitly mentioned, 2) the cell’s value is greater/less than the mentioned value in question, 3) the cell’s value is maximum/minimum over the whole column if the question involves superlative words.
- Passage Retriever: it aims to link cells implicitly mentioned by the question through its hyperlinked passage. The linking model consists of a TD-IDF retriever with 2-3 gram lexicon and a longest-substring retriever, this ensemble retriever calculates the distances with all the passages in the pool and highlight the ones with cosine distance lower than a threshold τ . The retrieved passages are mapped back to the linked cell in the table.

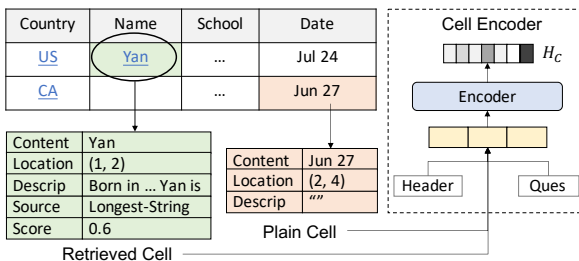


Figure 5: Illustration of cell encoder of retrieved (green) and plain cells (orange).

We call the set of cells from these two sources as “retrieved cells” denotes by \mathcal{C} . Each

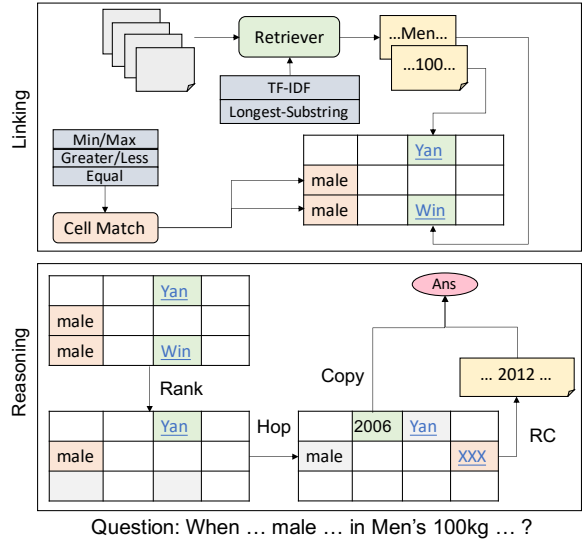


Figure 6: Illustration of the proposed model to perform multi-hop reasoning over table and passage.

retrieved cell c is encoded by 5-element tuple (content, location, description, source, score). Content represents the string representation in the table, Content refers to the absolute row and column index in the table, description refers to the evidence sentence in the hyperlinked passage, which gives highest similarity score to question, source denotes where the entry comes from (e.g. equal/argmax/passage/etc), score denotes the score of linked score normalized to $[0, 1]$.

Reasoning This phase is aimed to model the multi-hop reasoning in the table and passage, we specifically break down the whole process into three stages, namely the ranking stage $p_f(c|q, \mathcal{C})$, hoping stage $p_h(c'|q, c)$, and the reading comprehension stage $p_r(a|P, q)$. These three stages are modeled with three different neural networks. We first design a cell encoding scheme to encode each cell in the table as depicted in Figure 5: 1) for “retrieved cells”, it contains information for retrieval source and score, 2) for “plain cells” (not retrieved), we set the information in source and score to empty. We concatenate them with their table field and question, and then fed into a encoder module (BERT) to obtain its vector representation H_c .

1) Ranking model: As the “retriever cells“ contain many noises, we leverage a ranker model to predict the “correct” linked cells for the next stage. Specifically, this model takes each cell c along with its neighboring \mathcal{N}_c (cells in the same row) and

feed them all into the cell encoder to obtain their representations $\{H_c\}$. The representations are aggregated and further fed to a feed-forward neural network to obtain a score s_c , which is normalized over the whole set of linked cell \mathcal{C} as follows:

$$p_f(c|q, \mathcal{C}) = \frac{\exp(s_c)}{\sum_{c' \in \mathcal{C}} \exp(s_{c'})} \quad (1)$$

2) Hop model: this model takes the predicted cell from the previous stage and then decide which neighboring cell or itself to hop to. Specifically, we represent each hop pair $(c \rightarrow c')$ using their concatenated representation $H_{c,c'} = [H_c, H_{c'}]$. The representation is fed to a feed-forward neural network to obtain a hop score $s_{c,c'}$, which is normalized over all the possible end cells as follows:

$$p_f(c'|q, c) = \frac{\exp(s_{c,c'})}{\sum_{c'' \in \mathcal{N}_c \cup c} \exp(s_{c,c''})} \quad (2)$$

3) RC model: this model finally takes the hopped cell c from last stage and find answer from it. If the cell is not hyperlinked, the RC model will simply output its plain text as the answer, otherwise, the plain text of the cell is prepended to the linked passage $P(c)$ for reading comprehension. The prepended passage P and the question are given as the input to the question answering model to predict the score of answer's start and end index as $g_s(P, q, index)$ and $g_e(P, q, index)$, which are normalized over the whole passage $|P|$ to calculate the likelihood $p_r(a|P, q)$ as follows:

$$p_r(a|P, q) = \frac{\exp(g_s(P, q, a_s))}{\sum_{i \in |P|} \exp(g_s(P, q, i))} \frac{\exp(g_e(P, q, a_e))}{\sum_{i \in |P|} \exp(g_e(P, q, i))}$$

where a_s is the start index of answer a and a_e is the end index of answer a .

By breaking the reasoning process into three stages, we manage to cover the Type-I/II/III/VI questions well. For example, the Type-III question first uses the ranking model to select the most likely cell from retrievers, and then use the hop model to jump to neighboring hyperlinked cell, finally use the RC model to extract the answer.

Training & Inference The three-stage decomposition breaks the question answering likelihood $p(a|q, T)$ into the following marginal probability:

$$\sum_{c \in \mathcal{C}} p_f(c|q, \mathcal{C}) \sum_{c' \in \mathcal{N}_c; a \in P(c')} p_f(c'|c, q) p_r(a|P(c'), q)$$

where the marginalization is over all the linked cells c , and all the neighboring cell with answer a

in its plain text or linked passages. However, directly maximizing the marginal likelihood is unnecessarily complicated as the marginalization leads to huge computation cost. Therefore, we propose to train the three models independently and then combine them to do inference.

By using the source location of answers, we are able to 1) infer which cells c in the retrieved set \mathcal{C} are valid, which can be applied to train the ranking model, 2) infer which cell it hops to get the answer, which we can be applied to train the hop model. Though the synthesized reasoning paths are somewhat noisy, it is still enough to be used for training the separate models in a weakly supervised manner. For the RC model, we use the passages containing the ground-truth answer to train it. The independent training avoids the marginalization computation to greatly decrease the computation and time cost. During inference, we apply these three models sequentially to get the answer. Specifically, we use greedy search at first two steps to remain only the highest probably cell and finally extract the answer using the RC model.

5 Experiments

5.1 Experimental Setting

In the linking phase, we set the retrieval threshold τ to a specific value. All the passages having distance lower than τ will be retrieved and fed as input to the reasoning phase. If there is no passage that has been found with a distance lower than τ , we will simply use the document with the lowest distance as the retrieval result. Increasing τ can increase the recall of correct passages, but also increase the difficulty of the filter model in the reasoning step.

In the reasoning phase, we mainly utilize BERT (Devlin et al., 2019) as our encoder for the cells and passages due to its strong semantic understanding. Specifically, we use four BERT variants provided by huggingface library³, namely base-uncased, based-cased, large-uncased, and large-cased. We train the modules all for 3.0 epochs and save their checkpoint file at the end of each epoch. The filtering, hop, and RC models use AdamW (Loshchilov and Hutter, 2017) optimizer with learning rates of 2e-6, 5e-6, and 3e-5. We held out a small development set for model selection on the saved checkpoints and use the most performant ones in inference.

³<https://github.com/huggingface/transformers>

Model	Dev			Test		
	In-Table	In-Passage	Total	In-Table	In-Passage	Total
	EM/F1	EM/F1	EM/F1	EM/F1	EM/F1	EM/F1
Table-Only	14.7/19.1	2.4/4.5	8.4/12.1	14.2/18.8	2.6/4.7	8.3/11.7
Passage-Only	9.2/13.5	26.1/32.4	19.5/25.1	8.9/13.8	25.5/32.0	19.1/25.0
HYBRIDER (BERT-base-uncased, $\tau=0.7$)	51.2/58.6	39.6/46.4	42.9/50.0	50.9/58.6	37.4/45.7	41.8/49.5
HYBRIDER (BERT-base-uncased, $\tau=0.8$)	51.3/58.4	40.1/47.6	43.5/50.6	51.7/59.1	37.8/46.0	42.2/49.9
HYBRIDER (BERT-base-uncased, $\tau=0.9$)	51.5/58.6	40.5/47.9	43.7/50.9	52.1/59.3	38.1/46.3	42.5/50.2
HYBRIDER (BERT-large-uncased, $\tau=0.8$)	54.3/61.4	39.1/45.7	44.0/50.7	56.2/63.3	37.5/44.4	43.8/50.6

Table 5: Experimental results of different models, In-Table refers to the subset of questions which have their answers in the table, In-Passage refers to the subset of questions which have their answer in a certain passage.

5.2 Evaluation

Following previous work (Rajpurkar et al., 2016), we use exact match (EM) and F1 as two evaluation metrics. F1 metric measures the average overlap between the prediction and ground-truth answers. We assess human performance on a held-out set from the test set containing 500 instances. To evaluate human performance, we distribute each question along with its table to crowd-workers and compare their answer with the ground-truth answer. We obtain an estimated accuracy of EM=88.2 and F1=93.5, which is higher than both SQuAD (Rajpurkar et al., 2016) and HotpotQA (Yang et al., 2018). The higher accuracy is due to the In-Table questions (over 40%), which have much lesser ambiguity than the text-span questions.

5.3 Experimental Results

We demonstrate the experimental results for different models in Table 5, where we list fine-grained accuracy over the questions with answers in the cell and passage separately. The In-Table questions are remarkably simpler than In-Passage question because they do not the RC reasoning step; the overall accuracy is roughly 8-10% higher than its counterpart. With the experimented model variants, the best accuracy is achieved with BERT-large-uncased as backend, which can beat the BERT-base-uncased by roughly 2%. However, its performance is still far lagged behind human performance, leaving ample room for future research.

Heterogeneous Reasoning From Table 5, we can clearly observe that using either Table-Only or Passage-Only model achieves a poor accuracy below 20%. In contrast, the proposed HYBRIDER can achieve up to 50% EM increase by leveraging both structured and unstructured data during reasoning. It strongly supports the necessity to do heterogeneous reasoning in HYBRIDQA.

Retriever Threshold We also experiment with a different τ threshold. Having an aggressive retriever increases the recall of the mentioned cells, but it increases the burden for the ranking model. Having a passive retriever can guarantee the precision of predicted cells, but it also potentially miss evidence for the following reasoning phase. There exist trade-offs between these different modes. In Table 5, we experiment with different τ during the retrieval stage and find that the model is rather stable, which means the model is quite insensitive regarding different threshold values.

5.4 Error Analysis

To analyze the cause of the errors in HYBRIDER, we propose to break down into four types as Figure 7. Concretely, linking error is caused by the retriever/linker, which fails to retrieve the most relevant cell in the linking phase. In the reasoning phase: 1) ranking error is caused by the ranking model, which fails to assign a high score to the correct retrieved cell. 2) hop error occurs when the correctly ranked cell couldn't hop to the answer cell. 3) RC error refers to the hoped cell is correct, but the RC model fails to extract the correct text span from it. We perform our anal-




Figure 7: The error of HYBRIDER is based on its stages. Pink cell means the answer cell; green means the model's prediction; circle means the current cell.

ysis on the full dev set based on the bert-large-uncased model ($\tau=0.8$), as indicated in Figure 7, the errors are quite uniformly distributed into the four categories except the reading comprehension step is slightly more erroneous. Based on the

step-wise error, we can compute its product as $87.4\% \times 87.9\% \times 89.2\% \times 61.9\% \approx 42\%$ and find that the result consistent well the overall accuracy, which demonstrates the necessity to perform each reasoning step correctly. Such error cascading makes the problem extremely difficult than the previous homogeneous question answering problems.

By breaking down the reasoning into steps, HYBRIDER layouts strong explainability about its rationale, but it also causes error propagation, i.e., the mistakes made in the earlier stage are non-reversible in the following stage. We believe future research on building an end-to-end reasoning model could alleviate such an error propagation problem between different stages in HYBRIDER.

6 Related Work

Text-Based QA Since the surge of SQuAD (Rajpurkar et al., 2016) dataset, there have been numerous efforts to tackle the machine reading comprehension problem. Different datasets like DrQA (Chen et al., 2017), TriviaQA (Joshi et al., 2017), SearchQA (Dunn et al., 2017) and DROP (Dua et al., 2019) have been proposed. As the SQuAD (Rajpurkar et al., 2016) questions that are relatively simple because they usually require no more than one sentence in the paragraph to answer. The following datasets further challenge the QA model’s capability to handle different scenarios like open-domain, long context, multi-hop, discrete operations, etc. There has been a huge success in proving that the deep learning model can show strong competence in terms of understanding text-only evidence. Unlike these datasets, HYBRIDQA leverages structured information in the evidence form, where the existing models are not able to handle, which distinguishes it from the other datasets.

KB/Table-Based QA Structured knowledge is known as unambiguous and compositional, which absorbs lots of attention to the QA system built on KB/Tables. There have been multiple datasets like WebQuestion (Berant et al., 2013), ComplexWebQuestions (Talmor and Berant, 2018), WebQuestionSP (Yih et al., 2015) on using FreeBase to answer natural questions. Besides KB, structured or semi-structured tables are also an interesting form. Different datasets like WikiTableQuestions (Pasupat and Liang, 2015), WikiSQL (Zhong et al., 2017), SPIDER (Yu et al., 2018), TabFact (Chen et al., 2020) have been proposed to build models which can interact with such structured information.

However, both KB and tables are known to suffer from low coverage issues. Therefore, HYBRIDQA combine tables with text as complementary information to answer natural questions.

Hybrid QA There are some pioneering studies on designing hybrid question answering systems to aggregate heterogeneous information. GRAFT (Sun et al., 2018) proposes to use the early fusion system and use heuristics to build a question-specific subgraph that contains sentences from corpus and entities, facts from KB. PullNet (Sun et al., 2019) improves over GRAFT to use an integrated framework that dynamically learns to retrieve and reason over heterogeneous information to find the best answers. More recently, KARReader (Xiong et al., 2019) proposes to reformulate the questions in latent space by reading retrieved text snippets under KB-incomplete cases. These models simulate a ‘fake’ KB-incomplete scenario by masking out triples from KB. In contrast, the questions in HYBRIDQA are inherently hybrid in the sense that it requires both information forms to reason, which makes our testbed more realistic.

7 Conclusion

We present HYBRIDQA, which is collected as the first hybrid question answering dataset over both tabular and textual data. We release the data to facilitate the current research on using heterogeneous information to answer real-world questions. We design HYBRIDER as a strong baseline and offer interesting insights about the model. We believe HYBRIDQA is an interesting yet challenging next-problem for the community to solve.

Acknowledgement

The authors would like to thank the anonymous reviewers for their thoughtful comments.

References

- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1533–1544.
- Chandra Sekhar Bhagavatula, Thanapon Noraset, and Doug Downey. 2013. Methods for exploring and mining tables on wikipedia. In *Proceedings of the ACM SIGKDD Workshop on Interactive Data Exploration and Analytics*, pages 18–26.

- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879.
- Jifan Chen and Greg Durrett. 2019. Understanding dataset design choices for multi-hop reasoning. 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 4026–4032.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyong Zhou, and William Yang Wang. 2020. Tabfact: A large-scale dataset for table-based fact verification. *International Conference on Learning Representations (ICLR)*.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases. *arXiv preprint arXiv:1909.03683*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. 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 4171–4186.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs. 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 2368–2378.
- Matthew Dunn, Levent Sagun, Mike Higgins, V Ugur Guney, Volkan Cirik, and Kyunghyun Cho. 2017. Searchqa: A new q&a dataset augmented with context from a search engine. *arXiv preprint arXiv:1704.05179*.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611.
- Divyansh Kaushik and Zachary C Lipton. 2018. How much reading does reading comprehension require? a critical investigation of popular benchmarks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5010–5015.
- Tom Kwiatkowski, Eunsol Choi, Yoav Artzi, and Luke Zettlemoyer. 2013. Scaling semantic parsers with on-the-fly ontology matching. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1545–1556.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *ICLR 2019*.
- 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.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Haitian Sun, Tania Bedrax-Weiss, and William Cohen. 2019. Pullnet: Open domain question answering with iterative retrieval on knowledge bases and text. 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 2380–2390.
- Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William Cohen. 2018. Open domain question answering using early fusion of knowledge bases and text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4231–4242.
- Alon Talmor and Jonathan Berant. 2018. The web as a knowledge-base for answering complex questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 641–651.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.
- Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. 2019. Improving question answering over incomplete kbs with knowledge-aware reader. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4258–4264.

- Xiaojun Xu, Chang Liu, and Dawn Song. 2017. Sqlnet: Generating structured queries from natural language without reinforcement learning. *arXiv preprint arXiv:1711.04436*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380.
- Wen-tau Yih, Ming-Wei Chang, Xiaodong He, and Jianfeng Gao. 2015. Semantic parsing via staged query graph generation: Question answering with knowledge base. 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 1321–1331.
- Wen-tau Yih, Matthew Richardson, Christopher Meek, Ming-Wei Chang, and Jina Suh. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–206.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921.
- Yuyu Zhang, Hanjun Dai, Zornitsa Kozareva, Alexander J Smola, and Le Song. 2018. Variational reasoning for question answering with knowledge graph. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*.