

# TWEETQA: A Social Media Focused Question Answering Dataset

Wenhan Xiong<sup>†</sup>, Jiawei Wu<sup>†</sup>, Hong Wang<sup>†</sup>, Vivek Kulkarni<sup>†</sup>,  
Mo Yu<sup>\*</sup>, Shiyu Chang<sup>\*</sup>, Xiaoxiao Guo<sup>\*</sup>, William Yang Wang<sup>†</sup>

<sup>†</sup> University of California, Santa Barbara

<sup>\*</sup> IBM Research

{xwhan, william}@cs.ucsb.edu, yum@us.ibm.com, {shiyu.chang, xiaoxiao.guo}@ibm.com

## Abstract

With social media becoming increasingly popular on which lots of news and real-time events are reported, developing automated question answering systems is critical to the effectiveness of many applications that rely on real-time knowledge. While previous datasets have concentrated on question answering (QA) for formal text like news and Wikipedia, we present the first large-scale dataset for QA over social media data. To ensure that the tweets we collected are useful, we only gather tweets used by journalists to write news articles. We then ask human annotators to write questions and answers upon these tweets. Unlike other QA datasets like SQuAD in which the answers are extractive, we allow the answers to be abstractive. We show that two recently proposed neural models that perform well on formal texts are limited in their performance when applied to our dataset. In addition, even the fine-tuned BERT model is still lagging behind human performance with a large margin. Our results thus point to the need of improved QA systems targeting social media text.<sup>1</sup>

## 1 Introduction

Social media is now becoming an important real-time information source, especially during natural disasters and emergencies. It is now very common for traditional news media to frequently probe users and resort to social media platforms to obtain real-time developments of events. According to a recent survey by Pew Research Center<sup>2</sup>, in 2017, more than two-thirds of Americans read some of their news on social media. Even for American people who are 50 or older, 55% of them report getting news from social media,

<sup>1</sup>The Dataset can be found at <https://tweetqa.github.io/>.

<sup>2</sup><http://www.journalism.org/2017/09/07/news-use-across-social-media-platforms-2017/>

---

**Passage:** *Oh man just read about Paul Walkers death. So young. Ugggh makes me sick especially when it's caused by an accident. God bless his soul. – Jay Sean (@jaysean) December 1, 2013*

---

**Q:** *why is sean torn over the actor's death?*  
**A:** *walker was young*

---

Table 1: An example showing challenges of TWEETQA. Note the highly informal nature of the text and the presence of social media specific text like usernames which need to be comprehended to accurately answer the question.

which is 10% points higher than the number in 2016. Among all major social media sites, Twitter is most frequently used as a news source, with 74% of its users obtaining their news from Twitter. All these statistical facts suggest that understanding user-generated noisy social media text from Twitter is a significant task.

In recent years, while several tools for core natural language understanding tasks involving syntactic and semantic analysis have been developed for noisy social media text (Gimpel et al., 2011; Ritter et al., 2011; Kong et al., 2014; Wang et al., 2014), there is little work on question answering or reading comprehension over social media, with the primary bottleneck being the lack of available datasets. We observe that recently proposed QA datasets usually focus on formal domains, e.g. CNN/DAILYMAIL (Hermann et al., 2015) and NewsQA (Trischler et al., 2016) on news articles; SQuAD (Rajpurkar et al., 2016) and WIKI-MOVIES (Miller et al., 2016) that use Wikipedia.

In this paper, we propose the first large-scale dataset for QA over social media data. Rather than naively obtaining tweets from Twitter using

the Twitter API<sup>3</sup> which can yield irrelevant tweets with no valuable information, we restrict ourselves only to tweets which have been used by journalists in news articles thus implicitly implying that such tweets contain useful and relevant information. To obtain such relevant tweets, we crawled thousands of news articles that include tweet quotations and then employed crowd-sourcing to elicit questions and answers based on these event-aligned tweets. Table 1 gives an example from our TWEETQA dataset. It shows that QA over tweets raises challenges not only because of the informal nature of oral-style texts (*e.g.* inferring the answer from multiple short sentences, like the phrase “so young” that forms an independent sentence in the example), but also from tweet-specific expressions (such as inferring that it is “Jay Sean” feeling sad about Paul’s death because he posted the tweet).

Furthermore, we show the distinctive nature of TWEETQA by comparing the collected data with traditional QA datasets collected primarily from formal domains. In particular, we demonstrate empirically that three strong neural models which achieve good performance on formal data do not generalize well to social media data, bringing out challenges to developing QA systems that work well on social media domains.

In summary, our contributions are:

- We present the first question answering dataset, TWEETQA, that focuses on social media context;
- We conduct extensive analysis of questions and answer tuples derived from social media text and distinguish it from standard question answering datasets constructed from formal-text domains;
- Finally, we show the challenges of question answering on social media text by quantifying the performance gap between human readers and recently proposed neural models, and also provide insights on the difficulties by analyzing the decomposed performance over different question types.

## 2 Related Work

**Tweet NLP** Traditional core NLP research typically focuses on English newswire datasets such as the Penn Treebank (Marcus et al., 1993). In recent

years, with the increasing usage of social media platforms, several NLP techniques and datasets for processing social media text have been proposed. For example, Gimpel et al. (2011) build a Twitter part-of-speech tagger based on 1,827 manually annotated tweets. Ritter et al. (2011) annotated 800 tweets, and performed an empirical study for part-of-speech tagging and chunking on a new Twitter dataset. They also investigated the task of Twitter Named Entity Recognition, utilizing a dataset of 2,400 annotated tweets. Kong et al. (2014) annotated 929 tweets, and built the first dependency parser for tweets, whereas Wang et al. (2014) built the Chinese counterpart based on 1,000 annotated Weibo posts. To the best of our knowledge, question answering and reading comprehension over short and noisy social media data are rarely studied in NLP, and our annotated dataset is also an order of magnitude large than the above public social-media datasets.

**Reading Comprehension** Machine reading comprehension (RC) aims to answer questions by comprehending evidence from passages. This direction has recently drawn much attention due to the fast development of deep learning techniques and large-scale datasets. The early development of the RC datasets focuses on either the cloze-style (Hermann et al., 2015; Hill et al., 2015) or quiz-style problems (Richardson et al., 2013; Lai et al., 2017). The former one aims to generate single-token answers from automatically constructed pseudo-questions while the latter requires choosing from multiple answer candidates. However, such unnatural settings make them fail to serve as the standard QA benchmarks. Instead, researchers started to ask human annotators to create questions and answers given passages in a crowdsourced way. Such efforts give the rise of large-scale human-annotated RC datasets, many of which are quite popular in the community such as SQuAD (Rajpurkar et al., 2016), MS MARCO (Nguyen et al., 2016), NewsQA (Trischler et al., 2016). More recently, researchers propose even challenging datasets that require QA within dialogue or conversational context (Reddy et al., 2018; Choi et al., 2018). According to the difference of the answer format, these datasets can be further divided to two major categories: *extractive* and *abstractive*. In the first category, the answers are in text spans of the given passages, while in the latter case, the answers may

<sup>3</sup><https://developer.twitter.com/>

not appear in the passages. It is worth mentioning that in almost all previously developed datasets, the passages are from Wikipedia, news articles or fiction stories, which are considered as the formal language. Yet, there is little effort on RC over informal one like tweets.

### 3 TweetQA

In this section, we first describe the three-step data collection process of TWEETQA: *tweet crawling*, *question-answer writing* and *answer validation*. Next, we define the specific task of TWEETQA and discuss several evaluation metrics. To better understand the characteristics of the TWEETQA task, we also include our analysis on the answer and question characteristics using a subset of QA pairs from the development set.

#### 3.1 Data Collection

**Tweet Crawling** One major challenge of building a QA dataset on tweets is the sparsity of informative tweets. Many users write tweets to express their feelings or emotions about their personal lives. These tweets are generally uninformative and also very difficult to ask questions about. Given the linguistic variance of tweets, it is generally hard to directly distinguish those tweets from informative ones. In terms of this, rather than starting from Twitter API Search, we look into the archived snapshots<sup>4</sup> of two major news websites (CNN, NBC), and then extract the tweet blocks that are embedded in the news articles. In order to get enough data, we first extract the URLs of all section pages (*e.g.* World, Politics, Money, Tech) from the snapshot of each home page and then crawl all articles with tweets from these section pages. Note that another possible way to collect informative tweets is to download the tweets that are posted by the official Twitter accounts of news media. However, these tweets are often just the summaries of news articles, which are written in formal text. As our focus is to develop a dataset for QA on informal social media text, we do not consider this approach.

After we extracted tweets from archived news articles, we observed that there is still a portion of tweets that have very simple semantic structures and thus are very difficult to raise meaningful questions. An example of such tweets can be like:

<sup>4</sup><https://archive.org/>

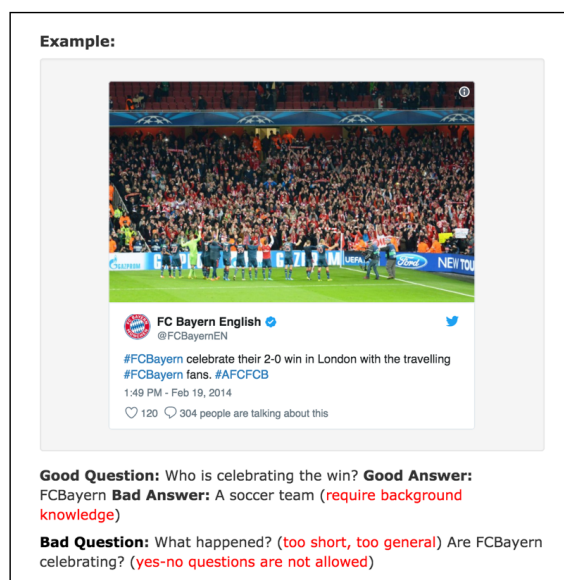


Figure 1: An example we use to guide the crowdworkers when eliciting question answer pairs. We elicit question that are neither too specific nor too general, do not require background knowledge.

“*Wanted to share this today - @IAMSteveHarvey*”. This tweet is actually talking about an image attached to this tweet. Some other tweets with simple text structures may talk about an inserted link or even videos. To filter out these tweets that heavily rely on attached media to convey information, we utilize a state-of-the-art semantic role labeling model trained on CoNLL-2005 (He et al., 2017) to analyze the predicate-argument structure of the tweets collected from news articles and keep only the tweets with more than two labeled arguments. This filtering process also automatically filters out most of the short tweets. For the tweets collected from CNN, 22.8% of them were filtered via semantic role labeling. For tweets from NBC, 24.1% of the tweets were filtered.

**Question-Answer Writing** We then use Amazon Mechanical Turk to collect question-answer pairs for the filtered tweets. For each Human Intelligence Task (HIT), we ask the worker to read three tweets and write two question-answer pairs for each tweet. To ensure the quality, we require the workers to be located in major English speaking countries (*i.e.* Canada, US, and UK) and have an acceptance rate larger than 95%. Since we use tweets as context, lots of important information are contained in hashtags or even emojis. Instead of only showing the text to the workers, we use *javascript* to directly embed the whole tweet into

each HIT. This gives workers the same experience as reading tweets via web browsers and help them to better compose questions.

To avoid trivial questions that can be simply answered by superficial text matching methods or too challenging questions that require background knowledge. We explicitly state the following items in the HIT instructions for question writing:

- No Yes-no questions should be asked.
- The question should have at least five words.
- Videos, images or inserted links should not be considered.
- No background knowledge should be required to answer the question.

To help the workers better follow the instructions, we also include a representative example showing both good and bad questions or answers in our instructions. Figure 1 shows the example we use to guide the workers.

As for the answers, since the context we consider is relatively shorter than the context of previous datasets, we do not restrict the answers to be in the tweet, otherwise, the task may potentially be simplified as a classification problem. The workers are allowed to write their answers in their own words. We just require the answers to be brief and can be directly inferred from the tweets.

After we retrieve the QA pairs from all HITs, we conduct further post-filtering to filter out the pairs from workers that obviously do not follow instructions. We remove QA pairs with yes/no answers. Questions with less than five words are also filtered out. This process filtered 13% of the QA pairs. The dataset now includes 10,898 articles, 17,794 tweets, and 13,757 crowdsourced question-answer pairs. The collected QA pairs will be directly available to the public, and we will provide a script to download the original tweets and detailed documentation on how we build our dataset. Also note that since we keep the original news article and news titles for each tweet, our dataset can also be used to explore more challenging generation tasks. Table 2 shows the statistics of our current collection, and the frequency of different types of questions is shown in Table 3. All QA pairs were written by 492 individual workers.

Dataset Statistics	
# of Training triples	10,692
# of Development triples	1,086
# of Test triples	1,979
Average question length (#words)	6.95
Average answer length (#words)	2.45

Table 2: Basic statistics of TWEETQA

Question Type	Percentage
What	42.33%
Who	29.36%
How	7.79%
Where	7.00%
Why	2.61%
Which	2.43%
When	2.16%
Others	6.32%

Table 3: Question Type statistics of TWEETQA

**Answer Validation** For the purposes of human performance evaluation and inter-annotator agreement checking, we launch a different set of HITs to ask workers to answer questions in the test and development set. The workers are shown with the tweet blocks as well as the questions collected in the previous step. At this step, workers are allowed to label the questions as “NA” if they think the questions are not answerable. We find that 3.1% of the questions are labeled as unanswerable by the workers (for SQuAD, the ratio is 2.6%). Since the answers collected at this step and previous step are written by different workers, the answers can be written in different text forms even they are semantically equal to each other. For example, one answer can be “*Hillary Clinton*” while the other is “*@HillaryClinton*”. As it is not straightforward to automatically calculate the overall agreement, we manually check the agreement on a subset of 200 random samples from the development set and ask an independent human moderator to verify the result. It turns out that 90% of the answers pairs are semantically equivalent, 2% of them are partially equivalent (one of them is incomplete) and 8% are totally inconsistent. The answers collected at this step are also used to measure the human performance. We have 59 individual workers participated in this process.

### 3.2 Task and Evaluation

As described in the *question-answer writing* process, the answers in our dataset are different from those in some existing extractive datasets.



Thus we consider the task of answer generation for TWEETQA and we use several standard metrics for natural language generation to evaluate QA systems on our dataset, namely we consider BLEU-1<sup>5</sup> (Papineni et al., 2002), Meteor (Denkowski and Lavie, 2011) and Rouge-L (Lin, 2004) in this paper.

To evaluate machine systems, we compute the scores using both the original answer and validation answer as references. For human performance, we use the validation answers as generated ones and the original answers as references to calculate the scores.

### 3.3 Analysis

In this section, we analyze our dataset and outline the key properties that distinguish it from standard QA datasets like SQuAD (Rajpurkar et al., 2016). First, our dataset is derived from social media text which can be quite informal and user-centric as opposed to SQuAD which is derived from Wikipedia and hence more formal in nature. We observe that the shared vocabulary between SQuAD and TWEETQA is only 10.79%, suggesting a significant difference in their lexical content. Figure 2 shows the 1000 most distinctive words in each domain as extracted from SQuAD and TWEETQA. Note the stark differences in the words seen in the TWEETQA dataset, which include a large number of user accounts with a heavy tail. Examples include @realdonaldtrump, @jdsutter, @justinkirkland and #cnworldcup, #goldenglobes. In contrast, the SQuAD dataset rarely has usernames or hashtags that are used to signify events or refer to the authors. It is also worth noting that the data collected from social media can not only capture events and developments in real-time but also capture individual opinions and thus requires reasoning related to the authorship of the content as is illustrated in Table 1. In addition, while SQuAD requires all answers to be spans from the given passage, we do not enforce any such restriction and answers can be free-form text. In fact, we observed that 43% of our QA pairs consists of answers which do not have an exact substring matching with their corresponding passages. All of the above distinguishing factors have implications to existing models

<sup>5</sup>The answer phrases in our dataset are relatively short so we do not consider other BLEU scores in our experiments

which we analyze in upcoming sections.

We conduct analysis on a subset of TWEETQA to get a better understanding of the kind of reasoning skills that are required to answer these questions. We sample 150 questions from the development set, then manually label their reasoning categories. Table 4 shows the analysis results. We use some of the categories in SQuAD (Rajpurkar et al., 2016) and also proposes some tweet-specific reasoning types.

Our first observation is that almost half of the questions only require the ability to identify phrases. Although most of the “paraphrasing only” questions are considered as fairly easy questions, we find that a significant amount (about 3/4) of these questions are asked about event-related topics, such as information about “who did what to whom, when and where”. This is actually consistent with our motivation to create TWEETQA, as we expect this dataset could be used to develop systems that automatically collect information about real-time events.

Apart from these questions, there are also a group of questions that require understanding common sense, deep semantics (*i.e.* the answers cannot be derived from the literal meanings of the tweets), and relations of sentences<sup>6</sup> (including coreference resolution), which are also appeared in other RC datasets (Rajpurkar et al., 2016). On the other hand, the TWEETQA also has its unique properties. Specifically, a significant amount of questions require certain reasoning skills that are specific to social media data:

- **Understanding authorship:** Since tweets are highly personal, it is critical to understand how questions/tweets related to the authors.
- **Oral English & Tweet English:** Tweets are often oral and informal. QA over tweets requires the understanding of common oral English. Our TWEETQA also requires understanding some tweet-specific English, like conversation-style English.
- **Understanding of user IDs & hashtags:** Tweets often contains user IDs and hashtags, which are single special tokens. Understanding these special tokens is important to answer person- or event-related questions.

<sup>6</sup>There are more instances of this reasoning type compared to formal datasets since tweets are usually short sentences.

Type	Fraction (%)	Example
Paraphrasing only	47.3	P: <i>Belgium camp is 32 miles from canceled game at US base. Surprised Klinsmann didn't offer to use his helicopter pilot skills to give a ride. – Grant Wahl (@GrantWahl)</i> Q: <i>what expertise does klinsmann possess?</i> A: <i>helicopter pilot skills</i>
<u>Types Beyond Paraphrasing</u>		
Sentence relations	10.7	P: <i>My heart is hurting. You were an amazing tv daddy! Proud and honored to have worked with one of the best. Love and Prayers #DavidCassidy— Alexa PenaVega (@alexavega) November 22, 2017</i> Q: <i>who was an amazing tv daddy?</i> A: <i>#davidcassidy</i>
<u>Authorship</u>	17.3	P: <i>Oh man just read about Paul Walkers death. So young. Ugggh makes me sick especially when it's caused by an accident. God bless his soul. – Jay Sean (@jaysean)</i> Q: <i>why is sean torn over the actor's death?</i> A: <i>walker was young</i>
<u>Oral/Tweet English habits</u>	10.7	P: <i>I got two ways to watch the OLYMPICS!! CHEAH!! USA!! Leslie Jones (@Lesdoggg) August 6, 2016</i> Q: <i>who is being cheered for?</i> A: <i>usa</i>
<u>UserIDs &amp; Hashtags</u>	12.0	P: <i>Started researching this novel in 2009. Now it is almost ready for you to read. Excited! #InTheUnlikelyEvent – Judy Blume (@judyblume)</i> Q: <i>what is the name of the novel?</i> A: <i>in the unlikely event.</i>
Other commonsense	6.7	P: <i>Don't have to be Sherlock Holmes to figure out what Russia is up to ... – Lindsey Graham (@LindseyGrahamSC)</i> Q: <i>what literary character is referenced?</i> A: <i>sherlock holmes.</i>
Deep semantic	3.3	P: <i>@MayorMark its all fun and games now wait until we are old enough to vote #lastlaugh – Dylan (@DFPFilms1)</i> Q: <i>when does the author suggest a change?</i> A: <i>when he's of voting age.</i>
Ambiguous (Meaningless questions)	5.3	P: <i>The #endangeredriver would be a sexy bastard in this channel if it had water. Quick turns. Narrow. (I'm losing it) – John D. Sutter (@jdsutter)</i> Q: <i>what is this user "losing"</i> A: <i>he is losing it</i>

Table 4: Types of reasoning abilities required by TWEETQA. Underline indicates tweet-specific reasoning types, which are common in TWEETQA but are rarely observed in previous QA datasets. Note that the first type represents questions that *only* require the ability of paraphrasing, while the rest of the types require some other more salient abilities besides paraphrasing. **Overlaps could exist between different reasoning types in the table.** For example, the *second* example requires both the understanding of sentences relations and tweet language habits to answer the question; and the *third* example requires both the understanding of sentences relations and authorship.



Figure 2: Visualization of vocabulary differences between SQuAD (left) and TWEETQA (right). Note the presence of a heavy tail of hash-tags and usernames on TWEETQA that are rarely found on SQuAD. The color range from red to gray indicates the frequency (red the highest and gray the lowest).

## 4 Experiments

To show the challenge of TweetQA for existing approaches, we consider four representative methods as baselines. For data processing, we first remove the URLs in the tweets and then tokenize the QA pairs and tweets using NLTK.<sup>7</sup> This process is consistent for all baselines.

### 4.1 Query Matching Baseline

We first consider a simple query matching baseline similar to the IR baseline in Kociský et al. (2017). But instead of only considering several genres of spans as potential answers, we try to match the question with all possible spans in the tweet context and choose the span with the highest BLEU-1 score as the final answer, which follows the method and implementation<sup>8</sup> of answer span selection for open-domain QA (Wang et al., 2017). We include this baseline to show that TWEETQA is a nontrivial task which cannot be easily solved with superficial text matching.

### 4.2 Neural Baselines

We then explore three typical neural models that perform well on existing formal-text datasets. One takes a generative perspective and learns to decode the answer conditioned on the question and context, while the others learn to extract a text span from the context that best answers the question.

**Generative QA** RNN-based encoder-decoder models (Cho et al., 2014; Bahdanau et al., 2014) have been widely used for natural language generation tasks. Here we consider a recently pro-

posed generative model (Song et al., 2017) that first encodes the context and question into a multi-perspective memory via four different neural matching layers, then decodes the answer using an attention-based model equipped with both copy and coverage mechanisms. The model is trained on our dataset for 15 epochs and we choose the model parameters that achieve the best BLEU-1 score on the development set.

**BiDAF** Unlike the aforementioned generative model, the Bi-Directional Attention Flow (BiDAF) (Seo et al., 2016) network learns to directly predict the answer span in the context. BiDAF first utilizes multi-level embedding layers to encode both the question and context, then uses bi-directional attention flow to get a query-aware context representation, which is further modeled by an RNN layer to make the span predictions. Since our TWEETQA does not have labeled answer spans as in SQuAD, we need to use the human-written answers to retrieve the answer-span labels for training. To get the approximate answer spans, we consider the same matching approach as in the query matching baseline. But instead of using questions to do matching, we use the human-written answers to get the spans that achieve the best BLEU-1 scores.

**Fine-Tuning BERT** This is another extractive RC model that benefits from the recent advance in pretrained general language encoders (Peters et al., 2018; Devlin et al., 2018). In our work, we select the BERT model (Devlin et al., 2018) which has achieved the best performance on SQuAD. In our experiments, we use the PyTorch reimple-

<sup>7</sup><http://www.nltk.org>

<sup>8</sup><https://github.com/shuohangwang/mprc>

Models	Evaluation on Dev/Test Data		
	BLEU-1	METEOR	ROUGE-L
HUMAN	76.4 78.2	63.7 66.7	70.9 73.5
EXTRACT-UB	79.5 80.3	68.8 69.8	74.3 75.6
Query-Matching	30.3 29.4	12.0 12.1	17.0 17.4
<i>Neural Baselines</i>			
BiDAF	48.3 48.7	31.6 31.4	38.9 38.6
Generative	53.4 53.7	32.1 31.8	39.5 39.0
BERT	67.3 69.6	56.9 58.6	62.6 64.1

Table 5: Overall performance of baseline models. EXTRACT-UB refers to our estimation of the upper bound of extractive methods.

mentation<sup>9</sup> of the uncased base model. The batch size is set as 12 and we fine-tune the model for 2 epochs with learning rate 3e-5.

## 5 Evaluation

### 5.1 Overall Performance

We test the performance of all baseline systems using the three generative metrics mentioned in Section 3.2. As shown in Table 5, there is a large performance gap between human performance and all baseline methods, including BERT, which has achieved superhuman performance on SQuAD. This confirms that TWEETQA is more challenging than formal-test RC tasks.

We also show the upper bound of the extractive models (denoted as EXTRACT-UPPER). In the upper bound method, the answers are defined as n-grams from the tweets that maximize the BLEU-1/METEOR/ROUGE-L compared to the annotated groundtruth. From the results, we can see that the BERT model still lags behind the upper bound significantly, showing great potential for future research. It is also interesting to see that the HUMAN performance is slightly worse compared to the upper bound. This indicates (1) the difficulty of our problem also exists for human-beings and (2) for the answer verification process, the workers tend to also extract texts from tweets as answers.

According to the comparison between the two non-pretraining baselines, our generative baseline yields better results than BiDAF. We believe this is largely due to the abstractive nature of our dataset, since the workers can sometimes write the answers using their own words.

<sup>9</sup><https://github.com/huggingface/pytorch-pretrained-BERT>

### 5.2 Performance Analysis over Human-Labeled Question Types

Reasoning Types	Generative BERT	
	METEOR	ROUGE-L
Paraphrasing	37.6 73.4	44.1 81.8
Sentence relations	34.0 46.1	42.2 51.1
Authorship	38.4 55.9	46.1 61.9
Oral/Tweet habits	37.2 50.3	40.7 51.0 <sup>†</sup>
UserIDs&Hashtags	3.8 <sup>◊</sup>  13.0 <sup>†</sup>	9.9 <sup>◊</sup>  16.2 <sup>†</sup>
Commonsense	20.1 63.5	33.1 67.1
Deep semantics	7.19 <sup>◊</sup>  7.1 <sup>†</sup>	13.4 <sup>◊</sup>  10.3 <sup>†</sup>
Ambiguous	4.1 <sup>◊</sup>  25.0 <sup>†</sup>	11.0 <sup>◊</sup>  67.1

Table 6: BiDAF’s and the Generative model’s performance on questions that require different types of reasoning. ◊ and † denote the three most difficult reasoning types for the Generative and the BERT models.

To better understand the difficulty of the TWEETQA task for current neural models, we analyze the decomposed model performance on the different kinds of questions that require different types of reasoning (we tested on the subset which has been used for the analysis in Table 4). Table 6 shows the results of the best performed non-pretraining and pretraining approach, *i.e.*, the generative QA baseline and the fine-tuned BERT. Our full comparison including the BiDAF performance and evaluation on more metrics can be found in Appendix A. Following previous RC research, we also include analysis on automatically-labeled question types in Appendix B.

As indicated by the results on METEOR and ROUGE-L (also indicated by a third metric, BLEU-1, as shown in Appendix A), both baselines perform worse on questions that require the understanding deep semantics and userID&hashtags. The former kind of questions also appear in other benchmarks and is known to be challenging for many current models. The second kind of questions is tweet-specific and is related to specific properties of social media data. Since both models are designed for formal-text passages and there is no special treatment for understanding user IDs and hashtags, the performance is severely limited on the questions requiring such reasoning abilities. We believe that good segmentation, disambiguation and linking tools developed by the social media community for processing the userIDs and hashtags will significantly help these question types.



**On non-pretraining model** Besides the easy questions requiring mainly paraphrasing skill, we also find that the questions requiring the understanding of authorship and oral/tweet English habits are not very difficult. We think this is due to the reason that, except for these tweet-specific tokens, the rest parts of the questions are rather simple, which may require only simple reasoning skill (e.g. paraphrasing).

**On pretraining model** Although BERT was demonstrated to be a powerful tool for reading comprehension, this is the first time a detailed analysis has been done on its reasoning skills. From the results, the huge improvement of BERT mainly comes from two types. The first is paraphrasing, which is not surprising because that a well pretrained language model is expected to be able to better encode sentences. Thus the derived embedding space could work better for sentence comparison. The second type is commonsense, which is consistent with the good performance of BERT (Devlin et al., 2018) on SWAG (Zellers et al., 2018). We believe that this provides further evidence about the connection between large-scaled deep neural language model and certain kinds of commonsense.

## 6 Conclusion

We present the first dataset for QA on social media data by leveraging news media and crowdsourcing. The proposed dataset informs us of the distinctiveness of social media from formal domains in the context of QA. Specifically, we find that QA on social media requires systems to comprehend social media specific linguistic patterns like informality, hashtags, usernames, and authorship. These distinguishing linguistic factors bring up important problems for the research of QA that currently focuses on formal text. We see our dataset as a first step towards enabling not only a deeper understanding of natural language in social media but also rich applications that can extract essential real-time knowledge from social media.

## References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.

Kyunghyun Cho, Bart van Merriënboer, Çağlar

Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*.

Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentaoh Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. Quac: Question answering in context. *arXiv preprint arXiv:1808.07036*.

Michael J. Denkowski and Alon Lavie. 2011. Meteor 1.3: Automatic metric for reliable optimization and evaluation of machine translation systems. In *WMT@EMNLP*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Kevin Gimpel, Nathan Schneider, Brendan T. O’Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. 2011. Part-of-speech tagging for twitter: Annotation, features, and experiments. In *ACL*.

Luheng He, Kenton Lee, Mike Lewis, and Luke S. Zettlemoyer. 2017. Deep semantic role labeling: What works and what’s next. In *ACL*.

Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proc. of Conf. on Advances in NIPS*.

Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2015. The goldilocks principle: Reading children’s books with explicit memory representations. *arXiv preprint arXiv:1511.02301*.

Tomás Kociský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2017. The narrativeqa reading comprehension challenge. *CoRR*, abs/1712.07040.

Lingpeng Kong, Nathan Schneider, Swabha Swayamdipta, Archana Bhatia, Chris Dyer, and Noah A. Smith. 2014. A dependency parser for tweets. In *EMNLP*.

Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. RACE: Large-scale reading comprehension dataset from examinations. *Proc. of Conf. on EMNLP*.

Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*.

Mitchell P Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The penn treebank. *Computational linguistics*, 19(2):313–330.

- Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. Key-value memory networks for directly reading documents. *EMNLP*.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*.
- Kishore Papineni, Salim E. Roucos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proc. of Conf. on EMNLP*.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2018. Coqa: A conversational question answering challenge. *arXiv preprint arXiv:1808.07042*.
- Matthew Richardson, Christopher JC Burges, and Erin Renshaw. 2013. MCTest: A challenge dataset for the open-domain machine comprehension of text. In *Proc. of Conf. on EMNLP*.
- Alan Ritter, Sam Clark, Oren Etzioni, et al. 2011. Named entity recognition in tweets: an experimental study. In *Proceedings of the conference on empirical methods in natural language processing*, pages 1524–1534. Association for Computational Linguistics.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *CoRR*, abs/1611.01603.
- Linfeng Song, Zhiguo Wang, and Wael Hamza. 2017. A unified query-based generative model for question generation and question answering. *CoRR*, abs/1709.01058.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2016. NewsQA: A machine comprehension dataset. *arXiv preprint arXiv:1611.09830*.
- Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerald Tesauro, Bowen Zhou, and Jing Jiang. 2017. R3: Reinforced reader-ranker for open-domain question answering. *arXiv preprint arXiv:1709.00023*.
- William Yang Wang, Lingpeng Kong, Kathryn Mazaitis, and William W. Cohen. 2014. Dependency parsing for weibo: An efficient probabilistic logic programming approach. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, Doha, Qatar. ACL.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. Swag: A large-scale adversarial dataset for grounded commonsense inference. *arXiv preprint arXiv:1808.05326*.

## **A Full results of Performance Analysis over Human-Labeled Question Types**

Table 7 gives our full evaluation on human annotated question types.

Compared with the BiDAF model, one interesting observation is that the generative baseline gets much worse results on ambiguous questions. We conjecture that although these questions are meaningless, they still have many words that overlapped with the contexts. This can give BiDAF potential advantage over the generative baseline.

## **B Performance Analysis over Automatically-Labeled Question Types**

Besides the analysis on different reasoning types, we also look into the performance over questions with different first tokens in the development set, which provide us an automatic categorization of questions. According to the results in Table 8, the three neural baselines all perform the best on “Who” and “Where” questions, to which the answers are often named entities. Since the tweet contexts are short, there are only a small number of named entities to choose from, which could make the answer pattern easy to learn. On the other hand, the neural models fail to perform well on the “Why” questions, and the results of neural baselines are even worse than that of the matching baseline. We find that these questions generally have longer answer phrases than other types of questions, with the average answer length being 3.74 compared to 2.13 for any other types. Also, since all the answers are written by humans instead of just spans from the context, these abstract answers can make it even harder for current models to handle. We also observe that when people write “Why” questions, they tend to copy word spans from the tweet, potentially making the task easier for the matching baseline.

Reasoning Types	BLEU-1	METEOR	ROUGE-L
	BiDAF Generative BERT		
Paraphrasing	49.1 56.8 81.7	35.4 37.6 73.4	44.5 44.1 81.8
Sentence relations	43.3 53.4 50.0	26.8 34.0 46.1	32.8 42.2 51.1
Authorship	52.5 65.4 63.0	30.5 38.4 55.9	42.3 46.1 61.9
Oral/Tweet habits	45.8 60.8 60.4	34.8 37.2 50.3	35.1 40.7 51.0 <sup>†</sup>
UserIDs&Hashtags	30.0* 41.5 <sup>◊</sup>  29.3 <sup>†</sup>	8.30* 3.81 <sup>◊</sup>  13.0 <sup>†</sup>	13.7* 9.88 <sup>◊</sup>  16.2 <sup>†</sup>
Commonsense	27.6* 38.1 <sup>◊</sup>  72.9	22.4* 20.1 63.5	31.0* 33.1 67.1
Deep semantics	34.8* 53.8 25.0 <sup>†</sup>	7.85* 7.19 <sup>◊</sup>  7.1 <sup>†</sup>	17.5* 13.4 <sup>◊</sup>  10.3 <sup>†</sup>
Ambiguous	35.1 18.1 <sup>◊</sup>  31.6 <sup>†</sup>	29.2 4.11 <sup>◊</sup>  25.0 <sup>†</sup>	34.3 11.0 <sup>◊</sup>  67.1

Table 7: BiDAF’s and the Generative model’s performance on questions that require different types of reasoning. \*, ◊ and † denote the three most difficult reasoning types for BiDAF/Generative/BERT models.

Models	First-Word Question Types							
	What	Who	How	Where	When	Why	Which	Others
HUMAN	74.1	83.5	61.1	74.8	72.2	66.0	76.8	76.0
Query-Matching	32.4	29.8	28.4	27.1	22.9	51.9	22.7	21.1
	<i>Neural Baselines</i>							
BiDAF	44.5	54.9	41.0	60.2	46.5	36.1	44.7	41.6
Generative	46.8	63.8	53.4	61.7	45.4	44.3	51.4	43.1
BERT	64.8	72.5	57.7	78.1	64.5	61.0	67.2	59.2

Table 8: BLEU-1 scores on different types of questions. Calculated on the development set.