

ComQA: A Community-sourced Dataset for Complex Factoid Question Answering with Paraphrase Clusters

Abdalghani Abujabal¹, Rishiraj Saha Roy¹, Mohamed Yahya² and Gerhard Weikum¹

¹Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

{abujabal, rishiraj, weikum}@mpi-inf.mpg.de

²Bloomberg L.P., London, United Kingdom

myahya6@bloomberg.net

Abstract

To bridge the gap between the capabilities of the state-of-the-art in factoid question answering (QA) and what real users ask, we need large datasets of real user questions that capture the various question phenomena users are interested in, and the diverse ways in which these questions are formulated. We introduce *ComQA*, a large dataset of real user questions that exhibit different *challenging aspects* such as *temporal reasoning*, *compositionality*, etc. ComQA questions come from the WikiAnswers community QA platform¹. Through a large crowdsourcing effort, we clean the question dataset, group questions into *paraphrase clusters*, and annotate clusters with their answers. ComQA contains 11,214 questions grouped into 4,834 paraphrase clusters. We detail the process of constructing ComQA, including the measures taken to ensure its high quality while making effective use of crowdsourcing. We also present an extensive analysis of the dataset and the results achieved by state-of-the-art systems on ComQA, demonstrating that our dataset can be a driver of future research on QA.

1 Introduction

Factoid QA is the task of answering questions whose answer is one or a small number of entities (Voorhees and Tice, 2000). To advance research in QA in a manner consistent with the needs of end users, it is important to have access to benchmarks that reflect *real* user information needs by covering various *question phenomena* and the wide *lexical and syntactic variety* in expressing these information needs. The benchmarks should be *large* enough to facilitate the use of data-hungry machine learning methods. In this paper, we present ComQA, a *large dataset*

¹Such questions are usually ones that existing search engines and QA systems could not satisfactorily answer.



Figure 1: Paraphrase clusters from ComQA, covering a range of question aspects, with lexical and syntactic diversity.

of 11,214 *real* user questions collected from the WikiAnswers community QA website. As shown in Figure 1, the dataset contains *various question phenomena*. ComQA questions are grouped into 4,834 *paraphrase clusters* through a large-scale crowdsourcing effort, which capture lexical and syntactic variety. Crowdsourcing is also used to pair paraphrase clusters with answers to serve as a supervision signal for training and as a basis for evaluation.

Table 1 contrasts ComQA with other publicly available QA datasets. The foremost issue ComQA tackles is ensuring research is driven by real information needs formulated by real users. Most large-scale benchmarks resort to highly-templatic synthetically generated natural language questions (Bordes et al., 2015; Cai and Yates, 2013; Su et al., 2016; Talmor and Berant, 2018; Trivedi et al., 2017). Other benchmarks utilize

Dataset	Large scale (> 5K)	Real Information Needs	Complex Questions	Question Paraphrases
ComQA (This paper)	✓	✓	✓	✓
Free917 (Cai and Yates, 2013)	✗	✗	✗	✗
WebQuestions (Berant et al., 2013)	✓	✓	✗	✗
SimpleQuestions (Bordes et al., 2015)	✓	✗	✗	✗
QALD (Usbeck et al., 2017)	✗	✗	✓	✗
LC-QuAD (Trivedi et al., 2017)	✓	✗	✓	✗
ComplexQuestions (Bao et al., 2016)	✗	✓	✓	✗
GraphQuestions (Su et al., 2016)	✓	✗	✓	✓
ComplexWebQuestions (Talmor and Berant, 2018)	✓	✗	✓	✗
TREC (Voorhees and Tice, 2000)	✗	✓	✓	✗

Table 1: Comparison of ComQA with existing QA datasets over various dimensions.

search engine logs to collect their questions (Berant et al., 2013), which creates a bias towards simpler questions that search engines can already answer reasonably well. In contrast, ComQA questions come from WikiAnswers, a community QA website where users pose questions to be answered by other users. This is often a reflection of the fact that such questions are beyond the capabilities of commercial search engines and QA systems. Questions in our dataset exhibit a wide range of interesting aspects such as the need for temporal reasoning (Figure 1, cluster 1), comparison (e.g., comparatives, superlatives, ordinals) (Figure 1, cluster 2), compositionality (multiple, possibly nested, subquestions with multiple entities) (Figure 1, cluster 3), and unanswerable questions e.g., Figure 1, cluster 4.

ComQA is the result of a carefully designed large-scale crowdsourcing effort to group questions into paraphrase clusters and pair them with answers. Past work has demonstrated the benefits of paraphrasing for QA (Abujabal et al., 2018; Berant and Liang, 2014; Dong et al., 2017; Fader et al., 2013). Motivated by this, we judiciously use crowdsourcing to obtain clean paraphrase clusters from WikiAnswer’s noisy ones, resulting in ones like those shown in Figure 1, with both lexical and syntactic variations. The only other existing dataset to provide such clusters is that of Su et al. (2016), but that is based on synthetic information needs.

For answering, recent research has shown that combining various resources for answering significantly improves performance (Savenkov and Agichtein, 2016; Sun et al., 2018; Xu et al., 2016). Therefore, unlike earlier work, we do not pair ComQA with a specific knowledge base (KB) or text corpus for answering. We call on the research community to innovate in combining different answering sources to tackle ComQA and advance re-

search in QA. We also use crowdsourcing to pair paraphrase clusters with answers. ComQA answers are primarily Wikipedia entity URIs. This has two motivations: (i) it builds on the example of search engines that use Wikipedia as a primary way of answering entity-centric queries (e.g., through knowledge cards), and (ii) most modern KBs ground their entities in Wikipedia. Wherever the answers are temporal or measurable quantities, TIMEX3² and the International System of Units (SI)³ are used for normalization. Providing canonical answers allows for better comparison of different systems.

We present an extensive analysis of ComQA, where we introduce the various question phenomena in the dataset. Finally, we analyze the results of running state-of-the-art QA systems on ComQA. The main result is that ComQA exposes major shortcomings in these systems, mainly related to their inability to handle compositionality, time, and comparison. Our detailed error analysis provides inspiration for avenues of future work to ensure that QA systems meet the expectations of real users. To summarize, we make the following contributions:

- We present a dataset of 11,214 real user questions collected from a community QA website. The questions exhibit a range of aspects that are important for users and challenging for existing QA systems. Using crowdsourcing, questions are grouped into 4,834 paraphrase clusters that are annotated with answers. ComQA is available at: <http://qa.mpi-inf.mpg.de/comqa>.
- We present an extensive analysis of the dataset, and quantify the various difficulties found within. We also present the results of state-of-the-art QA systems on ComQA, and a detailed error analysis.

²<http://www.timeml.org>

³<https://en.wikipedia.org/wiki/SI>

2 Related Work

There are two main variants of the factoid QA task, with the distinction tied to the underlying resources used for answering and the nature of these answers. Traditionally, the problem of QA has been explored over large textual corpora (Cui et al., 2005; Dietz and Gamari, 2017; Ferrucci, 2012; Harabagiu et al., 2001, 2003; Ravichandran and Hovy, 2002; Saquete et al., 2009; Voorhees and Tice, 2000) with answers being textual phrases. More recently the problem has been explored over large structured resources such as knowledge bases (Berant et al., 2013; Unger et al., 2012; Yahya et al., 2013), with answers being semantically grounded entities. Very recent work demonstrated that the two variants are complementary, and a combination of the two results in the best answering performance (Savenkov and Agichtein, 2016; Sun et al., 2018; Xu et al., 2016).

QA over textual corpora. QA has a long tradition in IR and NLP, including benchmarking tasks in TREC (Voorhees and Tice, 2000; Dietz and Gamari, 2017) and CLEF (Magnini et al., 2004; Herrera et al., 2004). This has predominantly focused on retrieving answers from textual sources (Ferrucci, 2012; Harabagiu et al., 2006; Prager et al., 2004; Ravichandran and Hovy, 2002; Saquete et al., 2004, 2009; Yin et al., 2015). In IBM Watson (Ferrucci, 2012), structured data played a role, but text was the main source for answers, combined with learned models for question types. The TREC QA evaluation series provide hundreds of questions to be answered over a collection of documents, which have become widely adopted benchmarks for answer sentence selection (Wang and Nyberg, 2015). ComQA is orders of magnitude larger than TREC QA.

Reading comprehension is a recently introduced task, where the goal is to answer a question from a *given* textual paragraph (Kociský et al., 2017; Lai et al., 2017; Rajpurkar et al., 2016; Trischler et al., 2017; Yang et al., 2015). This setting is different from factoid QA, where the goal is to answer questions from a large repository of data (be it textual or structured), and not a single paragraph.

QA over knowledge bases. Recent efforts have focused on natural language questions as an interface for KBs, where questions are translated to structured queries via semantic parsing (Bao et al., 2016; Bast and Haussmann, 2015; Berant et al., 2013; Fader et al., 2013; Reddy et al.,

2014; Mohammed et al., 2018; Xu et al., 2016; Yang et al., 2014; Yao and Durme, 2014; Yih et al., 2015). Over the past five years, many datasets were introduced for this setting. However, as Table 1 shows, they are either small in size (QALD, Free917, and ComplexQuestions), composed of synthetically generated questions (SimpleQuestions, GraphQuestions, LC-QuAD and ComplexWebQuestions), or are structurally simple (WebQuestions). ComQA tackles all these shortcomings. The ability to return semantic entities as answers allows users to further explore these entities in various resources such as their Wikipedia pages, Freebase entries, etc. It also allows QA systems to tap into various interlinked resources for improvement (e.g., to obtain better lexicons, or train better NER systems). Because of this, ComQA provides semantically grounded reference answers where possible. ComQA answers are primarily Wikipedia entities (without committing to Wikipedia as an answering resource). For numerics and dates, ComQA adopts the SI and TIMEX3 standards, respectively.

3 Overview

In this work, a factoid question is a question whose answer is one or a small number of entities or literal values (Voorhees and Tice, 2000). For example, “*Who were the secretaries of state under Barack Obama?*” and “*When was Germany’s first post-war chancellor born?*”.

3.1 Questions in ComQA

Questions: A question in our dataset can exhibit *one or more* of the following phenomena:

- **Simple:** These are questions that ask about a single property of a named entity. E.g.: “*Where was Einstein born?*”
- **Compositional:** A question is compositional if obtaining its answer requires answering more primitive questions and combining these. These can be *intersection* or *nested* questions. Intersection questions are ones where two or more subquestions can be answered independently, and their answers intersected (e.g., “*Which films featuring Tom Hanks did Spielberg direct?*”). Nested questions are those where the answer of one subquestion is necessary to answer another (“*Who were the parents of the thirteenth president of the US?*”).
- **Temporal:** These are questions that require

temporal reasoning for deriving the answer, be it explicit (e.g., ‘in 1998’), implicit (e.g., ‘during the WWI’), relative (e.g., ‘current’), or latent (e.g. ‘Who is the US president?’). Temporal questions also include those whose answer is an explicit temporal expression (“When did Trenton become New Jersey’s capital?”).

- **Comparison:** We consider three types of comparison questions, namely, *comparatives* (“Which rivers in Europe are longer than the Rhine?”), *superlatives* (“What is the population of the largest city in Egypt?”), and those containing *ordinals* (“What was the name of Elvis’s first movie?”).
- **Telegraphic** (Joshi et al., 2014): These are short questions formulated in an informal manner similar to keyword queries (“First president India?”). Systems that rely on linguistic analysis of questions often fail on such questions.
- **Answer tuple:** Where an answer is a tuple of connected entities as opposed to a single entity (“When and where did George H. Bush go to college, and what did he study?”).

3.2 Answers in ComQA

Recent work has showed that the choice of answering resource, or the combination of resources significantly improves answering performance (Savenkov and Agichtein, 2016; Sun et al., 2018; Xu et al., 2016). Inspired by this, ComQA is not tied to a specific resource for answering. To this end, answers in ComQA are Wikipedia URIs, wherever this is possible. This enables QA systems to combine different answering resources which are linked to Wikipedia (e.g., DBpedia, Freebase, Yago, Wikidata, Wikipedia, ClueWeb09-FACCC1, etc). This also enables seamless comparison across QA systems whose individual answering resources are different, but are linked to Wikipedia. Literal value answers follow the TIMEX3 and SI standards. An answer in ComQA can be:

- **Entity:** ComQA entities are grounded in Wikipedia. However, Wikipedia is inevitably incomplete, so answers that cannot be grounded in Wikipedia are represented as plain text. For example the answer for “What is the name of Kristen Stewart adopted brother?” is {Taylor Stewart, Dana Stewart}.
- **Literal value:** Temporal answers follow the

TIMEX3 standard. For measurable quantities, we follow the International System of Units.

- **Empty:** Some questions are based on false premises, and hence, are unanswerable e.g., “Who was the first human being on Mars?” (no human has been on Mars, yet). The correct answer to such questions is the empty set. Such questions allow systems to cope with these cases. Recent work has started looking at this problem (Rajpurkar et al., 2018).

4 Dataset Construction

Our goal is to collect a dataset of factoid questions that represent real information needs and cover a range of question phenomena. Moreover, we want to have different paraphrases for each question. To this end, we tap into the potential of community QA platforms. Questions posed there represent real information needs. Moreover, users of community QA platforms provide (noisy) annotations around questions. In this work, we exploit the annotations where users mark questions as duplicates as a basis for paraphrase clusters, and clean those. Concretely, we started with the WikiAnswers crawl by Fader et al. (2014). We obtained ComQA from this crawl primarily through a large scale crowdsourcing effort to ensure it is of high quality. We describe this effort in what follows.

The original resource curated by Fader et al. contains 763M questions. Questions in the crawl are grouped into 30M paraphrase clusters based on feedback from WikiAnswers users. This clustering has a low accuracy (Fader et al., 2014). Extracting factoid questions and cleaning the clusters are thus essential for a high-quality dataset.

4.1 Preprocessing of WikiAnswers

To remove non-factoid questions, we applied the following two filters: (i) removing questions starting with ‘why’, and (ii) removing questions containing words like *(dis)similarities*, *differences*, *(dis)advantages*, *benefits*, and their synonyms. Questions matching these filters require a narrative as an answer, and are therefore out of scope. We also removed questions with less than three or more than twenty words, as we found these to be typically noisy or non-factoid questions. This left us with about 21M questions belonging to 6.1M clusters.

To further focus on factoid questions, we automatically classified the remaining questions into

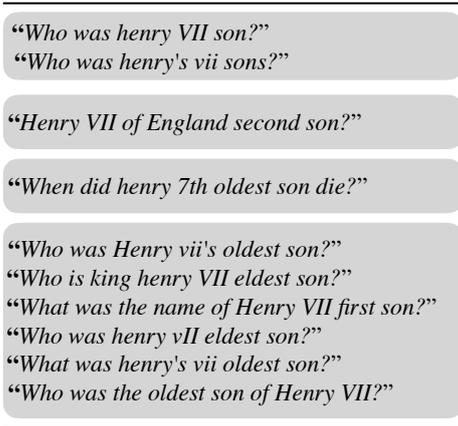


Figure 2: All ten questions belong to the same original WikiAnswers cluster. AMT Turkers split the original cluster into four new ones.

one or more of the following *four classes*: (1) temporal, (2) comparison, (3) single entity, and (4) multi-entity questions. We used SUTime (Chang and Manning, 2012) to identify temporal questions and the Stanford named entity recognizer (Finkel et al., 2005) to detect named entities. We used part-of-speech patterns to identify comparatives, superlatives, and ordinals. Clusters which did not have questions belonging to any of the above classes were discarded from further consideration. Although these clusters contain false negatives e.g., “*What official position did Mendeleev hold until his death?*” due to errors by the tagging tools, most discarded questions are out-of-scope e.g., “*How does the government help fight poverty?*”.

Manual inspection. We next applied the first stage of human curation to the dataset. Each WikiAnswers cluster was assigned to one of the four classes above based on the majority label of the questions within. We then randomly sampled 15K clusters from each of the four classes (60K clusters in total with 482K questions). We then sampled a representative question from each cluster at random (60K questions). We relied on the assumption that questions within the same cluster are semantically equivalent. These 60K questions were manually *examined by the authors* and those with unclear or non-factoid intent were removed along with the cluster that contains them. As a result, we ended up with 2.1K clusters with 13.7K questions.

4.2 Curating Paraphrase Clusters

We inspected a random subset of the 2.1K WikiAnswers clusters and found that questions in the same cluster are semantically *related but not equivalent*, which is line with observations in previous work (Fader et al., 2014). Dong et al. (2017) reported that 45% of question pairs were related rather than genuine paraphrases. For example, Figure 2 shows 10 questions in the same WikiAnswers cluster. Obtaining accurate paraphrase clusters is crucial to any systems that want to utilize them (Abujabal et al., 2018; Berant and Liang, 2014), and important to better understand our dataset. We therefore utilized crowdsourcing to clean the Wikianswers paraphrase clusters. We used Amazon Mechanical Turk (AMT) to identify semantically equivalent questions within a WikiAnswers cluster, thereby obtaining cleaner clusters for ComQA. Once we had the clean clusters, we set up a second AMT task to collect answers for each ComQA cluster of questions.

Task design. In designing the AMT task to clean up WikiAnswer’s paraphrase clusters, we had to ensure the simplicity of the task to obtain high quality results. Therefore, rather than giving workers a WikiAnswers cluster and asking them to partition it into clusters of paraphrases, we showed them pairs of questions from a cluster and asked them to make the binary decision of whether questions in the pair are paraphrases. To improve the efficiency of this annotation effort, we utilized the transitivity of the paraphrase relationship. Given a WikiAnswers cluster $Q = \{q_1, \dots, q_n\}$, we proceed in rounds to form ComQA paraphrase clusters. In the first round, we collect annotations for each pair (q_i, q_{i+1}) . The majority annotation among five annotators is taken. An initial clustering is formed accordingly, with clusters sharing the same question merged together (to account for transitivity). This process continues iteratively until no new clusters can be formed from a given WikiAnswers cluster.

Task statistics. We obtained annotations for 18,890 question pairs from 175 different workers. Each pair was shown to five different workers, with 65.7% of the pairs receiving unanimous agreement, 21.4% receiving four agreements and 12.9% receiving three agreements. By design, with five judges and binary annotations, no pair can have less three agreements. This resulted in questions being placed in paraphrase clusters, and

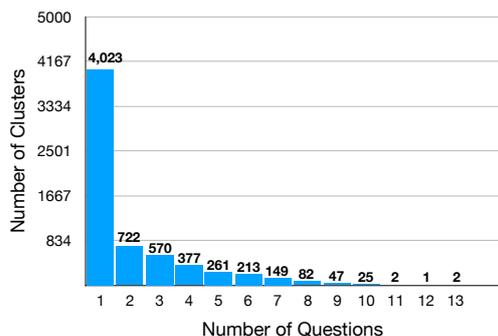


Figure 3: Distribution of the number of questions per cluster.

no questions were discarded at this stage. At the end of this step, the original 2.1K WikiAnswers clusters became 6.4K ComQA clusters with a total of 13.7K questions. Figure 3 shows the distribution of questions in clusters.

To test whether relying on the transitivity of the paraphrase relationship is suitable to reduce the annotation effort, we asked annotators to annotate 1,100 random pairs (q_1, q_3) , where we had already received positive annotations for the pairs (q_1, q_2) and (q_2, q_3) being paraphrases of each other. In 93.5% of the cases there was agreement. Additionally, as experts on the task, the authors manually assessed 600 pairs of questions, which serve as honeypots. There was 96.6% agreement with our annotations. An example result of this task is shown in Figure 2, where Turkers split the original WikiAnswers cluster into the four clusters shown.

4.3 Answering Questions

We were now in a position to obtain an answer annotation for each of the 6.4K clean clusters.

Task design. To collect answers, we designed another AMT task, where workers were shown a representative question randomly drawn from a cluster. Workers were asked to use the Web to find answers and to provide the URLs of Wikipedia entities that are suitable answers. Due to the inevitable incompleteness of Wikipedia, workers were asked to provide the surface form of an answer entity in case it does not have a Wikipedia page. If the answer is a full date, workers were asked to follow dd-mmm-yyyy format. For measurable quantities, workers were asked to provide units. We use TIMEX3 and the international system of units for normalizing temporal answers and measurable quantities e.g., ‘12th century’ to 11XX. If no answer is found, workers were asked to type

in ‘no answer’.

Task statistics. Each representative question was shown to three different workers. An answer is deemed correct if it is common between at least two workers. This resulted in 1.6K clusters (containing 2.4K questions) with no agreed-upon answers, which were dropped. We manually inspected some questions with no agreed-upon answers. Some questions were subjective, for example, “Who was the first democratically elected president of Mexico?”. Other questions received related answers e.g., “Who do the people in Iraq worship?” with Allah, Islam and Mohamed as answers from the three annotators. Other questions were underspecified e.g., “Who was elected the vice president in 1796?”, which misses the entity. At the end of the task, we ended up with 4,834 paraphrase clusters with 11,214 question-answer pairs, which form ComQA.

5 Dataset Analysis

In this section, we present a manual analysis of 300 questions sampled at random from the ComQA dataset. This analysis helps understand the different aspects of our dataset. A summary of the analysis is presented in Table 2.

Question categories. We categorized each question as either simple or complex. A question is deemed complex if it belongs to one or more of the compositional, temporal, or comparison classes (see Section 3). 56.33% of the questions were complex, with 32% compositional, 23.67% temporal, and 29.33% contain comparison conditions. Note that a question might contain a combination of conditions (“What country has the **highest** population in the year **2008**?” with comparison and temporal constraints).

We also identified questions of telegraphic nature e.g., “Julia Alvarez’s parents?”, with 8% of our questions being telegraphic. Such questions pose a challenge for systems that rely on linguistic analysis of questions (Joshi et al., 2014).

We counted the number of named entities in questions, with 23.67% containing two or more entities, indicating the compositional nature of questions. 2.67% have no entities e.g., “What public company has the most employees in the world?”. Such questions are hard since current methods assume the existence of a pivotal entity for each question.

Finally, 3.67% of questions are unanswerable,

Property	Example	Percentage%
<i>Compositional questions</i>		
Conjunction	“What is the capital of the country whose northern border is Poland and Germany?”	17.67
Nested	“When is Will Smith’s oldest son’s birthday?”	14.33
<i>Temporal questions</i>		
Explicit time	“Who was the winner of the World Series in 1994? ”	4.00
Implicit time	“Who was Britain’s leader during WWI? ”	4.00
Temporal answer	“ When did Trenton become New Jersey’s capital?”	15.67
<i>Comparison questions</i>		
Comparative	“Who was the first US president to serve 2 terms? ”	1.00
Superlative	“What ocean does the longest river in the world flow into?”	14.33
Ordinal	“When was Thomas Edisons first wife born?”	14.00
<i>Question formulation</i>		
wh- word	“ When did Trenton become New Jerseys capital?”	92.00
Telegraphic	“ Nevo first album?”	8.00
<i>Entity distribution in questions</i>		
Zero entity	“What public company has the most employees in the world?”	2.67
Single entity	“Who is Brad Walst ’s wife?”	75.67
Multi-entity	“What country in South America lies between Brazil and Argentina? ”	21.67
<i>Other features</i>		
Answer Tuple	“ Where was Peyton Manning born and what year was he born?”	2.00
Empty answer	“Who was Calgary’s first woman mayor?”	3.67

Table 2: Results of the manual analysis of 300 questions. Note that properties are not mutually exclusive.

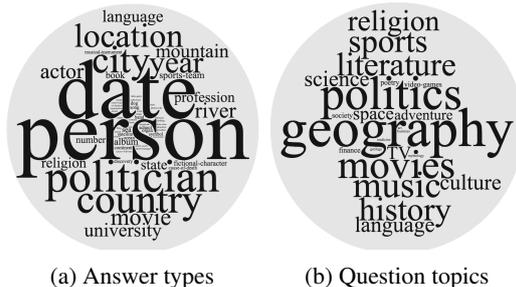


Figure 4: Answer types and question topics on 300 annotated examples as word clouds. The bigger the font, the more frequent the concept.

e.g., “Who was the first human being on Mars?”. Such questions incentivise QA systems to return non-empty answer sets only when suitable. We also compared ComQA with other current datasets over different question categories (Table 3).

Answer types. We annotated each question with the most fine-grained context-specific answer type (Ziegler et al., 2017). Answers in ComQA belong to a diverse set of types that range from coarse, e.g., *person* to fine, e.g., *sports manager*. Types also include literals e.g., *number* and *date*. Figure 4 (a) shows the set of answer types on 300 annotated examples as a word cloud.

Question topics. We annotated questions with topics they belong to: e.g., *geography*, *movies*, or *sports*. These are shown in Figure 4 (b), and

demonstrate the topical diversity of ComQA.

Question length. Questions in ComQA are fairly long, with a mean length of 7.73 words, indicating the compositional nature of questions.

6 Experiments

We experimented with various state-of-the-art QA systems, which achieved humble performance on ComQA, highlighting the need for new methods to handle the question phenomena within ComQA.

6.1 Experimental Setup

Splits. We generated a random split of 70% (7,850), 10% (1,121) and 20% (2,243), which serve as train, development and test sets.

Evaluation Metrics. We follow standard evaluation metrics adopted by the community: we compute average precision, recall, and F1 scores across all test questions. However, because ComQA includes unanswerable questions whose correct answer is the empty set, we define precision and recall to be 1 for a system that returns an empty set in response to an unanswerable question, and 0 otherwise (Rajpurkar et al., 2018).

6.2 Baselines

We evaluated two categories of QA systems that differ in the underlying answering resource: either KBs or textual extractions. We ran the follow-

Dataset	Size	Compositional	Temporal	Comparison	Telegraphic	Empty Answer
ComQA	11, 214	32%	24%	30%	8%	4%
WebQuestions (Berant et al., 2013)	5, 810	2%	7%	2%	0%	0%
ComplexQuestions (Bao et al., 2016)	2, 100	39%	34%	9%	0%	0%

Table 3: Comparison of ComQA with existing datasets over various phenomena. We manually annotated 100 random questions from each dataset.

	Avg. Prec	Avg. Rec	Avg. F1
Abujabal et al. (2017)	21.2	38.4	22.4
Bast and Haussmann (2015)	20.7	37.6	21.6
Berant and Liang (2015)	10.7	15.4	10.6
Berant et al. (2013)	13.7	20.1	12.0
Fader et al. (2013)	7.22	6.59	6.73

Table 4: Results of baselines on ComQA test set.

	WebQuestions F1	Free917 Accuracy	ComQA F1
Abujabal et al. (2017)	51.0	78.6	22.4
Bast and Haussmann (2015)	49.4	76.4	21.6
Berant and Liang (2015)	49.7	—	10.6
Berant et al. (2013)	35.7	62.0	12.0

Table 5: Results of baselines on different datasets.

ing systems using their publicly available code: (i) Abujabal et al. (2017), which automatically generates templates using question-answer pairs; (ii) Bast and Haussmann (2015), which instantiates query templates followed by query ranking; (iii) Berant and Liang (2015), which relies on agenda-based parsing and imitation learning; (iv) Berant et al. (2013), which uses rules to build queries from questions; and (v) Fader et al. (2013), which maps questions to queries over open vocabulary facts extracted from a large text corpus of Web documents. Note that our intention is not to assess the quality of current systems, but to show the challenging nature of ComQA.

All systems were run over the data sources for which they were designed. The first four baselines are over Freebase, therefore, we mapped ComQA answers (Wikipedia entities) to the corresponding Freebase names using the information stored with entities in Freebase. We observe that the Wikipedia answer entities have no counterpart in Freebase for only 7% of the ComQA questions. This suggests an oracle F1 score of 93.0. For Fader et al. (2013), which is over web extractions, we mapped Wikipedia URLs to their titles.

6.3 Results

Table 4 shows the performance of the baselines on the ComQA test set. Overall, the systems achieved

poor performance, suggesting that current methods cannot handle the complexity of our dataset, and that new models for QA are needed. Table 5 compares the performance of the systems on different datasets. For example, while Abujabal et al. (2017) achieved an F1 score of 51.0 on WebQuestions, it achieved 22.4 on ComQA.

The performance of Fader et al. (2013) is worse than the others due to the incompleteness of its underlying extractions and the complexity of ComQA questions that require higher-order relations and reasoning. However, the system answered complex questions, which KB-QA systems failed to answer. For example, it successfully answered “*What is the highest mountain in the state of Washington?*”. The answer to such a question is more readily available in Web text, compared to a KB, where more sophisticated reasoning would be required to handle the superlative. While text can easily answer questions like the one above, a slightly modified question e.g., “*What is the **fifth** highest mountain in the state of Washington?*” might not be explicitly found in text, which can be answered using KBs. Both examples above demonstrate the benefits of combining text and structured resources.

6.4 Error Analysis

For the two best performing systems on ComQA, QUINT (Abujabal et al., 2017) and AQQU (Bast and Haussmann, 2015), we manually inspected ComQA questions on which they failed, 100 questions per system. We classified failure sources into four categories: compositionality, temporal, comparison or NER. Table 6 shows the distribution of failure sources of both systems.

Compositionality. Both systems could not handle the compositional nature of questions. For example, they returned the father of Julius Caesar as an answer for “*What did Julius Caesar’s father work as?*”. However, the question requires another KB predicate that connects the father to his profession. For “*John Travolta and Jamie*

Lee Curtis starred in this movie?”, both systems returned the movies Jamie Lee Curtis appeared in, ignoring the part constraint that John Travolta should appear in them as well. Answering multi-relation questions over KBs remains an open problem given the large number of KB relations.

Temporal. Our analysis reveals that the tested systems often fail to capture temporal constraints in questions, be it explicit or implicit. For example, for *“Who won the Oscar for Best Actress in 1986?”*, both systems returned all winners and ignored the explicit year in the question together with the temporal aspect ‘in 1986’. Implicit temporal constraints e.g., named events like ‘Vietnam war’ in *“Who was the president of the US during Vietnam war?”* pose a challenge to current methods. Such constraints need first to be detected and normalized to a canonical time interval (November 1st, 1955 to April 30th, 1975 for the Vietnam war). Then, these systems need to compare the terms of the US presidents with above interval to account for the temporal relation of ‘during’. While detecting explicit time expressions can be done reasonably well using existing time taggers (Chang and Manning, 2012), identifying implicit ones is still a challenge. Furthermore, retrieving the correct temporal scopes of entities in questions (e.g., the terms of the US presidents) is hard due to the large number of temporal KB predicates associated with entities.

Comparison. Both systems perform poorly on comparison questions (comparatives, superlatives, and ordinals), which is expected since they were not designed to address those. To the best of our knowledge, no existing KB-QA system can handle comparison questions. Note that our goal is not to assess the quality of current methods, but to highlight that these methods miss categories of questions that are important to real users. For *“What is the first film Julie Andrews made?”* and *“What is the largest city in the state of Washington?”*, both systems returned the list of Julie Andrews’s films and the list of Washington’s cities, for the first and the second questions, respectively. While the first question requires the temporal attribute of `filmReleasedIn` to order by, the second question needs the attribute of `hasArea`. Identifying the correct attribute to order by as well as determining the order direction (ascending for the first and descending for the second) is challenging and out of scope for current methods.

Category	QUINT	AQQU
Compositionality error	39%	43%
Missing comparison	31%	26%
Missing temporal constraint	19%	22%
NER error	11%	9%

Table 6: Distribution of failure sources on ComQA questions on which QUINT and AQQU failed.

NER. NER errors come from false negatives, where entities are not detected, or false positives, where systems identify entities not intended as such. For example, in *“On what date did the Mexican Revolution end?”* QUINT identified ‘Mexican’, rather than ‘Mexican Revolution’ as an entity. For the latter case, the question *“What is the first real movie that was produced in 1903?”* does not ask about a specific entity. QUINT could not generate SPARQL queries and returned empty answer. Existing QA methods expect a pivotal entity in a question, which is not always the case.

Note that while baseline systems achieved low precision, they achieved higher recall (21.2 vs 38.4 for QUINT, respectively) (Table 4). This reflects the fact that these systems often cannot cope with the full complexity of ComQA questions, and instead end up evaluating underconstrained interpretations of the question.

To conclude, current methods can handle simple questions very well, but struggle with complex questions that involve multiple conditions on different entities or need to join the results from sub-questions. Handling such complex questions, however, is important if we are to satisfy information needs expressed by real users.

7 Conclusion

We presented ComQA, a large-scale dataset for factoid QA that harnesses a community QA platform, reflecting questions asked by real users. ComQA contains 11,214 questions paired with their answers, with questions grouped into paraphrase clusters through crowdsourcing. Questions exhibit a range of phenomena that state-of-the-art systems struggle with, as we demonstrated. ComQA is a challenging dataset that should help drive future research on factoid QA to match the needs of real users.

References

- Abdalghani Abujabal, Rishiraj Saha Roy, Mohamed Yahya, and Gerhard Weikum. 2018. Never-ending learning for open-domain question answering over knowledge bases. In *WWW*, pages 1053–1062.
- Abdalghani Abujabal, Mohamed Yahya, Mirek Riedewald, and Gerhard Weikum. 2017. Automated template generation for question answering over knowledge graphs. In *WWW*, pages 1191–1200.
- Jun-Wei Bao, Nan Duan, Zhao Yan, Ming Zhou, and Tiejun Zhao. 2016. Constraint-based question answering with knowledge graph. In *COLING*, pages 2503–2514.
- Hannah Bast and Elmar Haussmann. 2015. More accurate question answering on freebase. In *CIKM*, pages 1431–1440.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on Freebase from question-answer pairs. In *EMNLP*, pages 1533–1544.
- Jonathan Berant and Percy Liang. 2014. Semantic parsing via paraphrasing. In *ACL*, pages 1415–1425.
- Jonathan Berant and Percy Liang. 2015. Imitation learning of agenda-based semantic parsers. *TACL*, 3:545–558.
- Antoine Bordes, Nicolas Usunier, Sumit Chopra, and Jason Weston. 2015. Large-scale simple question answering with memory networks. *arXiv*.
- Qingqing Cai and Alexander Yates. 2013. Large-scale semantic parsing via schema matching and lexicon extension. In *ACL*, pages 423–433.
- Angel X. Chang and Christopher D. Manning. 2012. SUTIME: A library for recognizing and normalizing time expressions. In *LREC*, pages 3735–3740.
- Hang Cui, Renxu Sun, Keya Li, Min-Yen Kan, and Tat-Seng Chua. 2005. Question answering passage retrieval using dependency relations. In *SIGIR*, pages 400–407.
- Laura Dietz and Ben Gamari. 2017. TREC CAR: A data set for complex answer retrieval. version 1.4.(2017).
- Li Dong, Jonathan Mallinson, Siva Reddy, and Mirella Lapata. 2017. Learning to paraphrase for question answering. In *EMNLP*, pages 875–886.
- Anthony Fader, Luke Zettlemoyer, and Oren Etzioni. 2014. Open question answering over curated and extracted knowledge bases. In *KDD*, pages 1156–1165.
- Anthony Fader, Luke S. Zettlemoyer, and Oren Etzioni. 2013. Paraphrase-driven learning for open question answering. In *ACL*, pages 1608–1618.
- David A. Ferrucci. 2012. This is watson. *IBM Journal of Research and Development*, 56(3):1.
- Jenny Rose Finkel, Trond Grenager, and Christopher D. Manning. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *ACL*, pages 363–370.
- Sanda M. Harabagiu, V. Finley Lacatusu, and Andrew Hickl. 2006. Answering complex questions with random walk models. In *SIGIR*, pages 220–227.
- Sanda M. Harabagiu, Steven J. Maiorano, and Marius Pasca. 2003. Open-domain textual question answering techniques. *Natural Language Engineering*, 9(3):231–267.
- Sanda M. Harabagiu, Dan I. Moldovan, Marius Pasca, Rada Mihalcea, Mihai Surdeanu, Razvan C. Bunescu, Roxana Girju, Vasile Rus, and Paul Morarescu. 2001. The role of lexico-semantic feedback in open-domain textual question-answering. In *ACL*, pages 274–281.
- Jesús Herrera, Anselmo Peñas, and Felisa Verdejo. 2004. Question answering pilot task at CLEF 2004. In *CLEF*, pages 581–590.
- Mandar Joshi, Uma Sawant, and Soumen Chakrabarti. 2014. Knowledge graph and corpus driven segmentation and answer inference for telegraphic entity-seeking queries. In *EMNLP*, pages 1104–1114.
- 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.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard H. Hovy. 2017. RACE: large-scale reading comprehension dataset from examinations. In *EMNLP*, pages 785–794.
- Bernardo Magnini, Alessandro Vallin, Christelle Ayache, Gregor Erbach, Anselmo Peñas, Maarten de Rijke, Paulo Rocha, Kiril Ivanov Simov, and Richard F. E. Sutcliffe. 2004. Overview of the CLEF 2004 multilingual question answering track. In *CLEF*, pages 371–391.
- Salman Mohammed, Peng Shi, and Jimmy Lin. 2018. Strong baselines for simple question answering over knowledge graphs with and without neural networks. In *NAACL-HLT*, pages 291–296.
- John M. Prager, Jennifer Chu-Carroll, and Krzysztof Czuba. 2004. Question answering using constraint satisfaction: Qa-by-dossier-with-constraints. In *ACL*, pages 574–581.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *ACL*, pages 784–789.

- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100, 000+ questions for machine comprehension of text. In *EMNLP*, pages 2383–2392.
- Deepak Ravichandran and Eduard H. Hovy. 2002. Learning surface text patterns for a question answering system. In *ACL*, pages 41–47.
- Siva Reddy, Mirella Lapata, and Mark Steedman. 2014. Large-scale semantic parsing without question-answer pairs. *TACL*, pages 377–392.
- Estela Saquete, José Luis Vicedo González, Patricio Martínez-Barco, Rafael Muñoz, and Hector Llorens. 2009. Enhancing QA systems with complex temporal question processing capabilities. *J. Artif. Intell. Res.*
- Estela Saquete, Patricio Martínez-Barco, Rafael Muñoz, and José Luis Vicedo González. 2004. Splitting complex temporal questions for question answering systems. In *ACL*, pages 566–573.
- Denis Savenkov and Eugene Agichtein. 2016. When a Knowledge Base Is Not Enough: Question Answering over Knowledge Bases with External Text Data. In *SIGIR*, pages 235–244.
- Yu Su, Huan Sun, Brian Sadler, Mudhakar Srivatsa, Izzeddin Gur, Zenghui Yan, and Xifeng Yan. 2016. On generating characteristic-rich question sets for QA evaluation. In *EMNLP*, pages 562–572.
- Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William W Cohen. 2018. Open domain question answering using early fusion of knowledge bases and text. *EMNLP*.
- Alon Talmor and Jonathan Berant. 2018. The web as a knowledge-base for answering complex questions. In *NAACL*, pages 641–651.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. Newsqa: A machine comprehension dataset. In *Rep4NLP@ACL*, pages 191–200.
- Priyansh Trivedi, Gaurav Maheshwari, Mohnish Dubey, and Jens Lehmann. 2017. Lc-quad: A corpus for complex question answering over knowledge graphs. In *ISWC*, pages 210–218.
- Christina Unger, Lorenz Bühmann, Jens Lehmann, Axel-Cyrille Ngonga Ngomo, Daniel Gerber, and Philipp Cimiano. 2012. Template-based question answering over RDF data. In *WWW*, pages 639–648.
- Ricardo Usbeck, Axel-Cyrille Ngonga Ngomo, Bastian Haarmann, Anastasia Krithara, Michael Röder, and Giulio Napolitano. 2017. 7th Open Challenge on Question Answering over Linked Data (QALD-7). In *SemWebEval*.
- Ellen M. Voorhees and Dawn M. Tice. 2000. Building a question answering test collection. In *SIGIR*, pages 200–207.
- Di Wang and Eric Nyberg. 2015. A long short-term memory model for answer sentence selection in question answering. In *ACL*, pages 707–712.
- Kun Xu, Siva Reddy, Yansong Feng, Songfang Huang, and Dongyan Zhao. 2016. Question answering on Freebase via relation extraction and textual evidence. In *ACL*.
- Mohamed Yahya, Klaus Berberich, Shady Elbassuoni, and Gerhard Weikum. 2013. Robust question answering over the web of linked data. In *CIKM*, pages 1107–1116.
- Min-Chul Yang, Nan Duan, Ming Zhou, and Hae-Chang Rim. 2014. Joint relational embeddings for knowledge-based question answering. In *EMNLP*, pages 645–650.
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. Wikiqa: A challenge dataset for open-domain question answering. In *EMNLP*, pages 2013–2018.
- Xuchen Yao and Benjamin Van Durme. 2014. Information Extraction over Structured Data: Question Answering with Freebase. In *ACL*, pages 956–966.
- 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 *ACL*, pages 1321–1331.
- Pengcheng Yin, Nan Duan, Ben Kao, Junwei Bao, and Ming Zhou. 2015. Answering questions with complex semantic constraints on open knowledge bases. In *CIKM*, pages 1301–1310.
- David Ziegler, Abdalghani Abujabal, Rishiraj Saha Roy, and Gerhard Weikum. 2017. Efficiency-aware answering of compositional questions using answer type prediction. In *IJCNLP*, pages 222–227.