

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

Abdalghani Abujabal<sup>1</sup>, Rishiraj Saha Roy<sup>2</sup>, Mohamed Yahya<sup>3</sup> and Gerhard Weikum<sup>2</sup>

<sup>1</sup>Amazon Alexa, Aachen, Germany  
abujabaa@amazon.de

<sup>2</sup>Max Planck Institute for Informatics, Saarland Informatics Campus, Germany  
{rishiraj, weikum}@mpi-inf.mpg.de

<sup>3</sup>Bloomberg L.P., London, United Kingdom  
yahya.mohamed@gmail.com

## Abstract

To bridge the gap between the capabilities of the state-of-the-art in factoid question answering (QA) and what 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 *compositionality*, *temporal reasoning*, and *comparisons*. ComQA questions come from the WikiAnswers community QA platform, which typically contains questions that are not satisfactorily answerable by existing search engine technology. 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 natural language 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 datasets that reflect *real* user information needs by covering various *question phenomena* and the wide *lexical and syntactic variety* in expressing these information needs. The

<sup>1</sup>The main part of this work was carried out when the author was at the Max Planck Institute for Informatics.



Figure 1: ComQA paraphrase clusters covering a range of question aspects e.g., temporal and compositional questions, with lexical and syntactic diversity.

benchmarks should be large enough to facilitate the use of data-hungry machine learning methods. In this paper, we present ComQA, a large dataset 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 publicly available QA datasets. The foremost issue that ComQA tackles is ensuring research is driven by information needs formulated by real users. Most large-scale datasets resort to highly-templatic synthetically generated natural language questions (Bordes et al., 2015; Cai and Yates, 2013; Su et al.,

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.

2016; Talmor and Berant, 2018; Trivedi et al., 2017). Other datasets utilize 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 (Figure 1, cluster 2), compositionality (multiple subquestions with multiple entities and relations) (Figure 1, cluster 3), and unanswerable questions (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 WikiAnswers’ noisy ones, resulting in ones like those shown in Figure 1, with both lexical and syntactic variations. The only other 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, 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 research in QA. We use crowdsourcing to pair paraphrase clus-

ters with answers. ComQA answers are primarily Wikipedia entity URLs. This has two motivations: (i) it builds on the example of search engines that use Wikipedia entities as answers for 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, we use TIMEX3<sup>1</sup> and the International System of Units<sup>2</sup> 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 aspects of the dataset. We also analyze the results of running state-of-the-art QA systems on ComQA. 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, in this paper 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 and quantify the various difficulties in ComQA. We also present the results of state-of-the-art QA systems on ComQA, and a detailed error analysis.

<sup>1</sup><http://www.timeml.org>

<sup>2</sup><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 answering resources and the nature of answers. Traditionally, QA has been explored over large textual corpora (Cui et al., 2005; Harabagiu et al., 2001, 2003; Ravichandran and Hovy, 2002; Saquete et al., 2009) with answers being textual phrases. Recently, it has been explored over large structured resources such as KBs (Berant et al., 2013; Unger et al., 2012), with answers being semantic entities. Recent work demonstrated that the two variants are complementary, and a combination of the two results in the best performance (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; Saquete et al., 2004; Yin et al., 2015). In IBM Watson (Ferrucci, 2012), structured data played a role, but text was the main source for answers. The TREC QA evaluation series provide hundreds of questions to be answered over 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* (RC) 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. A recent direction in RC is dealing with unanswerable questions from the underlying data (Rajpurkar et al., 2018). ComQA includes such questions to allow tackling the same problem in the context of factoid QA.

**QA over KBs.** 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 Hausmann, 2015; Fader et al., 2013; Mohammed et al., 2018; Reddy et al., 2014; Yang et al., 2014; Yao and Durme, 2014; Yahya

et al., 2013). Over the past five years, many datasets were introduced for this setting. However, as Table 1 shows, they are either small in size (Free917, and ComplexQuestions), composed of synthetically generated questions (SimpleQuestions, GraphQuestions, LC-QuAD and ComplexWebQuestions), or are structurally simple (WebQuestions). ComQA addresses these shortcomings. Returning 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 in Wikipedia (without committing to Wikipedia as an answering resource). For numerical quantities and dates, ComQA adopts the International System of Units 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) e.g., “*Who were the secretaries of state under Barack Obama?*” and “*When was Germany’s first post-war chancellor born?*”.

### 3.1 Questions in ComQA

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

- **Simple:** questions about a single property of an entity (e.g., “*Where was Einstein born?*”)
- **Compositional:** A question is compositional if answering it 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?*”). In nested questions, the answer to 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: *comparatives* (“*Which rivers in Europe are longer than the Rhine?*”), *superlatives* (“*What is the population of the largest city in Egypt?*”), and *ordinal* questions (“*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 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 shown that the choice of answering resource, or the combination of resources significantly affects 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 primarily Wikipedia URLs. This enables QA systems to combine different answering resources which are linked to Wikipedia (e.g., DBpedia, Freebase, YAGO, Wikidata, etc). This also allows seamless comparison across these QA systems. 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:** In the factoid setting, some questions can be based on a false premise, 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 factoid questions that represent real information needs and cover a range of question aspects. 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 those platforms provide (noisy) annotations around questions e.g., paraphrase clusters. 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, which we describe 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 filtered out questions that (i) start with ‘*why*’, or (ii) contain words like (*dis*)*similarities*, *differences*, (*dis*)*advantages*, etc. Questions matching these filters are out of scope as they require a narrative answer. 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 questions into 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,

“Who was henry VII son?”  
 “Who was henry's vii sons?”

“Henry VII of England second son?”

“When did henry 7th oldest son die?”

“Who was Henry vii's oldest son?”  
 “Who is king henry VII eldest son?”  
 “What was the name of Henry VII first son?”  
 “Who was henry vii eldest son?”  
 “What was henry's vii oldest son?”  
 “Who was the oldest son of Henry VII?”

Figure 2: A WikiAnswers cluster split into four clusters by AMT Turkers.

most discarded questions are out-of-scope.

**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) and sampled a representative question from each of these clusters 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. We thus 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 necessarily equivalent*, which is in 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; Dong et al., 2017). 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




Figure 3: The distribution of questions in clusters.

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.

**Task design.** 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 the two questions are paraphrases. To reduce potentially redundant annotations, 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 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  $Q$ .

**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 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

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 <b>in 1994?</b> ”	4.00
Implicit time	“Who was Britain’s leader <b>during WWI?</b> ”	4.00
Temporal answer	“ <b>When</b> did Trenton become New Jersey’s capital?”	15.67
<i>Comparison questions</i>		
Comparative	“Who was the first US president to serve <b>2 terms?</b> ”	1.00
Superlative	“What ocean does the <b>longest</b> river in the world flow into?”	14.33
Ordinal	“When was Thomas Edisons <b>first</b> wife born?”	14.00
<i>Question formulation</i>		
Telegraphic	“Neyo 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 <b>Brad Walst</b> ’s wife?”	75.67
Multi-entity	“What country in <b>South America</b> lies between <b>Brazil</b> and <b>Argentina?</b> ”	21.67
<i>Other features</i>		
Answer tuple	“ <b>Where</b> was Peyton Manning born and <b>what year</b> 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.

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 corresponding URLs of Wikipedia entities. Due to the inevitable incompleteness of Wikipedia, workers were asked to provide the surface form of an answer entity if 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. For example, “Who was the first democratically elected president of Mexico?” is subjective. 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?”. At the end of the task, we ended up with 4,834 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 complex if it belongs to one or more of the compositional, temporal, or comparison classes. 56.33% of the questions were complex; 32% compositional, 23.67% temporal, and 29.33% contain

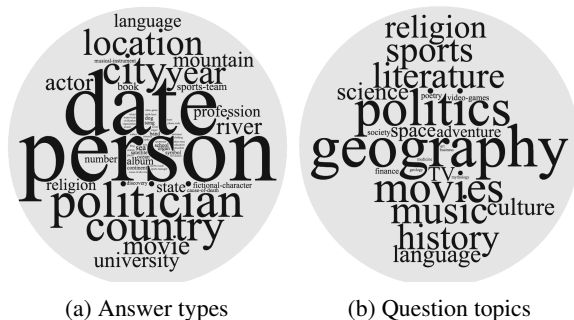


Figure 4: Answer types and question topics on 300 annotated examples as word clouds.

comparison conditions. A question may contain multiple conditions (“*What country has the **highest** population in the year **2008**?*” with comparison and temporal conditions).

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: 23.67% contain two or more entities, reflecting their compositional nature, and 2.67% contain no entities e.g., “*What public company has the most employees in the world?*”. Such questions can be hard as many methods assume the existence of a pivot entity in a question.

Finally, 3.67% of the questions are unanswerable, e.g., “*Who was the first human being on Mars?*”. Such questions incentivise QA systems to return non-empty answers only when suitable. In Table 3 we compare ComQA with other current datasets based on real user information needs over different question categories.

**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 such as `number` and `date`. Figure 4(a) shows answer types of the 300 annotated examples as a word cloud.

**Question topics.** We annotated questions with topics to which they belong (e.g., `geography`, `movies`, `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

In this section we present experimental results for running ComQA through state-of-the-art QA systems. Our experiments show that these systems achieve humble performance on ComQA. Through a detailed analysis, this performance can be attributed to systematic shortcomings in handling various question aspects in ComQA.

### 6.1 Experimental Setup

**Splits.** We partition ComQA into a random train/dev/test split of 70/10/20% with 7,850, 1,121 and 2,243 questions, respectively.

**Metrics.** We follow the community’s standard evaluation metrics: we compute average precision, recall, and F1 scores across all test questions. For 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, 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 following systems: (i) Abujabal et al. (2017), which automatically generates templates using question-answer pairs; (ii) Bast and Haussmann (2015), which instantiates hand-crafted 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 Web documents. Note that our intention is not to assess the quality of these systems, but to assess how challenging ComQA is.

The systems were trained with ComQA data. All systems were run over the data sources for which they were designed. The first four baselines are over Freebase. We therefore 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 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.

Dataset	Size	Compositional	Temporal	Comparison	Telegraphic	Empty Answer
<b>ComQA</b>	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 Hausmann (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 Hausmann (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.

### 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 (Free917 uses accuracy as a quality metric). 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 some complex questions, which KB-QA systems failed to answer. For example, it 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 is required to handle the superlative. However, a slightly modified question such as “*What is the **fourth** highest mountain in the state of Washington?*” is unlikely to be found in text, but be answered using KBs with the appropriate reasoning. 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 Hausmann, 2015), we manually inspected 100 questions on which they failed. We classified failure sources into four categories: compositionality, temporal, comparison or NER. Table 6 shows the distribution of these failure sources.

**Compositionality.** Neither system could 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?*”, while, 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 movies with Jamie Lee Curtis, ignoring the constraint that John Travolta should also appear in them. Properly answering multi-relation questions over KBs remains an open problem.

**Temporal.** Our analysis reveals that both systems fail to capture temporal constraints in questions, be it explicit or implicit. For “*Who won the Oscar for Best Actress in 1986?*”, they returned all winners and ignored the temporal restriction from ‘*in 1986*’. Implicit temporal constraints like named events (e.g., ‘*Vietnam war*’ in “*Who was the president of the US during Vietnam war?*”) pose a challenge to current methods. Such constraints need to be detected first and normalized to a canonical time interval (November 1st, 1955 to April 30th, 1975, for the Vietnam war). Then, systems need to compare the terms of the US presidents with the 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 difficult. Furthermore, retrieving the correct temporal scopes of entities in questions (e.g., the terms of the US presidents) is hard due to the large num-



ber of temporal KB predicates associated with entities.

**Comparison.** Both systems perform poorly on comparison questions, 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 attribute of `filmReleasedIn` to order by, the second 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.

**NER.** NER errors come from false negatives, where entities are not detected. For example, in “*On what date did the Mexican Revolution end?*” QUINT identified ‘*Mexican*’ rather than ‘*Mexican Revolution*’ as an entity. For “*What is the first real movie that was produced in 1903?*”, which does not ask about a specific entity, QUINT could not generate SPARQL queries. 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 dataset for QA that harnesses a community QA platform, reflecting ques-

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.

tions asked by real users. ComQA contains 11,214 question-answer pairs, with questions grouped into paraphrase clusters through crowdsourcing. Questions exhibit different aspects that current QA systems struggle with. ComQA is a challenging dataset that is aimed at driving future research on QA, to match the needs of real users.

## Acknowledgments

We would like to thank Tommaso Pasini for his helpful feedback.

## 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 Ayaiche, 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.
- 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.