

Look before you Hop: Conversational Question Answering over Knowledge Graphs Using Judicious Context Expansion

Philipp Christmann
MPI for Informatics, Germany
pchristm@mmci.uni-saarland.de

Rishiraj Saha Roy
MPI for Informatics, Germany
rishiraj@mpi-inf.mpg.de

Abdalghani Abujabal
Amazon Alexa, Germany
abujabaa@amazon.de

Jyotsna Singh
MPI for Informatics, Germany
jsingh@mpi-inf.mpg.de

Gerhard Weikum
MPI for Informatics, Germany
weikum@mpi-inf.mpg.de

ABSTRACT

Fact-centric information needs are rarely one-shot; users typically ask follow-up questions to explore a topic. In such a conversational setting, the user’s inputs are often incomplete, with entities or predicates left out, and ungrammatical phrases. This poses a huge challenge to question answering (QA) systems that typically rely on cues in full-fledged interrogative sentences. As a solution, we develop CONVEX: an unsupervised method that can answer incomplete questions over a knowledge graph (KG) by maintaining conversation context using entities and predicates seen so far and automatically inferring missing or ambiguous pieces for follow-up questions. The core of our method is a graph exploration algorithm that judiciously expands a frontier to find candidate answers for the current question. To evaluate CONVEX, we release CONVQUESTIONS, a crowdsourced benchmark with 11, 200 distinct conversations from five different domains. We show that CONVEX: (i) adds conversational support to any stand-alone QA system, and (ii) outperforms state-of-the-art baselines and question completion strategies.

ACM Reference Format:

Philipp Christmann, Rishiraj Saha Roy, Abdalghani Abujabal, Jyotsna Singh, and Gerhard Weikum. 2019. Look before you Hop: Conversational Question Answering over Knowledge Graphs Using Judicious Context Expansion. In *Proceedings of ACM CIKM Conference (CIKM’19)*. ACM, New York, NY, USA, Article 4, 10 pages. <https://doi.org/10.1145/nmmnnnnn.nmmnnnnn>

1 INTRODUCTION

1.1 Motivation

Obtaining direct answers to fact-centric questions is supported by large knowledge graphs (KGs) such as Wikidata or industrial KGs (at Google, Microsoft, Baidu, Amazon, etc.), consisting of semantically organized entities, attributes, and relations in the form of subject-predicate-object (SPO) triples. This task of question answering over KGs (KG-QA) has been intensively researched [2, 4, 5, 7, 15, 35, 36, 38]. However, users’ information needs are

not always expressed in well-formed and self-contained questions for one-shot processing. Quite often, users issue a series of follow-up questions to explore a topic [12, 30], analogous to search sessions [28]. A major challenge in such *conversational QA* settings is that follow-up questions are often incomplete, with entities or predicates not spelled out, and use of ungrammatical phrases. So a large part of the context is unspecified, assuming that the systems implicitly understand the user’s intent from previous interactions. Consider the following conversation as a running example. A user asks questions (or utterances) q^i and the system has to generate answers a^i :

q^0 : Which actor voiced the Unicorn in *The Last Unicorn*?
 a^0 : Mia Farrow
 q^1 : And Alan Arkin was behind . . . ?
 a^1 : Schmendrick
 q^2 : Who did the score?
 a^2 : Jimmy Webb
 q^3 : So who performed the songs?
 a^3 : America
 q^4 : Genre of this band’s music?
 a^4 : Folk rock, Soft rock
 q^5 : By the way, who was the director?
 a^5 : Jules Bass

Such conversations are characterized by a well-formed and complete initial question (q^0) with incomplete follow-ups ($q^1 - q^5$), an initial and often central entity of interest (“*The Last Unicorn*”), slight shifts in focus (inquiry of the band America’s genre in q^4), informal styles (q^1, q^5), and a running context comprised of entities and predicates in all preceding questions and answers (not just immediate precedents).

Limitations of state-of-the-art KG-QA. State-of-the-art systems [2, 7, 15, 21, 35] expect well-formed input questions (like q^0), complete with cue words for entities (“*Unicorn*”), predicates (“*voiced*”), and types (“*actor*”), and map them to corresponding KG-items. A SPARQL query (or an equivalent logical expression) is generated to retrieve answers. For example, a Wikidata query for q^0 would be: `SELECT ?x WHERE {TheLastUnicorn voiceActor ?x . ?x characterRole TheUnicorn}`. In our conversational setup, such methods completely fall apart due to the incompleteness of follow-up questions, and the ad-hoc ways in which they are phrased.

The alternative approach of *question completion* [17] aims to create syntactically correct full-fledged interrogative sentences from the user’s inputs, closing the gaps by learning from supervision

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM’19, October 2019, Beijing, China

© 2019 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM. . . \$15.00

<https://doi.org/10.1145/nmmnnnnn.nmmnnnnn>

pairs, while being agnostic to the underlying KG. However, this paradigm is bound to be limited and would fail for ad-hoc styles of user inputs or when training data is too sparse.

1.2 Approach and Contributions

Our proposed approach, CONVEX (CONVersational KG-QA with context EXpansion) overcomes these limitations, based on the following key ideas. The initial question is used to identify a small subgraph of the KG for retrieving answers, similar to what prior methods for unsupervised KG-QA use [7]. For incomplete and ungrammatical follow-up questions, we capture context in the form of a subgraph as well, and we dynamically maintain it as the conversation proceeds. This way, relevant entities and predicates from previous turns are kept in the gradually expanding context. However, we need to be careful about growing the subgraph too much as the conversation branches and broadens in scope. As nodes in a KG have many 1-hop neighbors and a huge number of 2-hop neighbors, there is a high risk of combinatorial explosion, and a huge subgraph would no longer focus on the topic of interest. CONVEX copes with this critical issue by judiciously expanding the context subgraph, using a combination of look-ahead, weighting, and pruning techniques. Hence the “look before you hop” in the paper title.

Specifically, CONVEX works as follows. Answers to the first question are obtained by any standard KG-QA system (we use the state-of-the-art system QAnswer [7] and other variants in our experiments over Wikidata). Entities in the initial question q^0 , the answer a^0 , and their connections initialize a context subgraph X^1 (X^t for turn t) for the conversation in the KG. When a follow-up question q^1 arrives, all nodes (entities, predicates, or types) in the KG-neighborhood of X^1 are deemed as candidates that will be used to expand the current graph. Brute force addition of all neighbors to X^1 will quickly lead to an explosion in its size after a few turns (hugely exacerbated if popular entities are added, e.g. Germany and Barcelona have $\approx 1.6M$ and $\approx 40k$ neighbor entities in Wikidata). Thus, we opt for prudent expansion as follows. Each neighbor is scored based on its similarity to the question, its distance to important nodes in X^1 , the conversation turn t , and KG priors. This information is stored in respective sorted lists with these neighbors as elements.

A small number of top-scoring neighbors of X^1 in a turn, termed “frontier nodes” ($\{F^1\}$), are identified by aggregating information across these queues. Next, all KG triples (SPO facts) for these frontiers *only*, are added to X^1 , to produce an expanded context X^1_+ . These $\{F^1\}$ are the most relevant nodes w.r.t the current question q^1 , and hence are expected to contain the answer a^1 in their close proximity. Each entity in X^1_+ is thus scored by its distance to each frontier node F^1_i and other important nodes in X^1 , and the top-ranked entity (possibly multiple, in case of ties) is returned as a^1 . This process is then iterated for each turn in the conversation with question q^t , producing $X^t, X^t_+, \{F^t\}$, and ultimately a^t at each step.

Benchmark. We compiled the first realistic benchmark, termed CONVQUESTIONS, for conversational KG-QA. It contains about 11k conversations which can be evaluated over Wikidata. They are compiled from the inputs of crowdworkers on Amazon Mechanical Turk, with conversations from five domains: “Books”, “Movies”, “Soccer”, “Music”, and “TV Series”. The questions feature a variety of complex question phenomena like comparisons, aggregations,

Notation	Concept
K, E, P, C, L	Knowledge graph, entity, predicate, class, literal
S, P, O	Subject, predicate, object
N, \mathcal{E}	Nodes and edges in graph
C, t	Conversation, turn
q^t, a^t	Question and answer at turn t
X^t, X^t_+	Initial and expanded context graphs at turn t
$N(X^t)$	k -hop neighborhood of nodes in X^t
$F^t = \{F^t_1 \dots F^t_r\}$	Frontier nodes at turn t
$E(q^t)$	Entities mapped to by words in q^t

Table 1: Notation for key concepts in CONVEX.

compositionality, and temporal reasoning. Answers are grounded in Wikidata entities to enable fair comparison across diverse methods.

Contributions. The main contributions of this work are:

- We devise CONVEX, an unsupervised method for addressing conversational question answering over knowledge graphs.
- We release CONVQUESTIONS, the first realistic benchmark to evaluate conversational KG-QA.
- We present extensive experiments, showing how CONVEX enables any stand-alone system with conversational support.
- An online demo and all code, data and results is available at <http://qa.mpi-inf.mpg.de/convex/>.

2 CONCEPTS AND NOTATION

We first introduce concepts that will assist in an easier explanation for the CONVEX method, and corresponding notations. An example workflow instantiating these concepts is shown in Fig. 1, and Table 1 provides a ready reference.

Knowledge graph. A knowledge graph, or a knowledge base, K is a set of subject-predicate-object (SPO) RDF triples, each representing a real-world fact, where S is of type entity E (like The Last Unicorn), P is a predicate (like director), and O is another entity, a class C (like animated feature film), or a literal L (like 19 November 1982). All E, P, C , and L in K are canonicalized. Most modern KGs support n -ary facts like movie-cast information (with more than two E and more than one P) via reification with intermediate nodes [32]. In Wikidata, such information is represented via optional qualifiers with the main fact (TheLastUnicorn voiceActor MiaFarrow . characterRole TheUnicorn). Compound Value Types (CVTs) were the Freebase analogue. Tapping into qualifier information is a challenge for SPARQL queries, but is easily accessible in a graph-based method like CONVEX.

CONVEX stores the KG as a graph $K = (N, \mathcal{E})$, with a set of nodes N and a set of edges \mathcal{E} , instead of a database-like RDF triple store. Each E, P, C , and L is assigned a unique node in K , with two nodes $n_1, n_2 \in N$ having an edge $e \in \mathcal{E}$ between them if there is a triple $\langle n_1, n_2, \cdot \rangle \in K$ or $\langle \cdot, n_1, n_2 \rangle \in K$. While it is more standard practice to treat each P as an edge label, we represent every item in K as a node, because it facilitates computing standard graph measures downstream. Examples of sample n and e are shown in the (sub-)graphs in Fig. 1. E and P nodes are in rectangles with sharp and rounded corners, respectively. For simplicity, C and L nodes are not shown. An important thing to note is that each instance of some P retains an *individual existence* in the graph to prevent false inferences (e.g. two voice actor nodes in the figure). As a




Figure 1: A typical conversation illustrating perfect (but simplified) context expansion and answering at every turn.

simple example, if we merge the node for married from two triples $\langle E_1, \text{married}, E_2 \rangle$ and $\langle E_3, \text{married}, E_4 \rangle$, then we may accidentally infer that E_1 is married to E_4 during answering.

Conversation. A conversation C with T turns is made up of a sequence of questions $Q = \{q^t\}$ and corresponding answers $A = \{a^t\}$, where $t = 0, 1, \dots, T$, such that $C = \langle (q^0, a^0), (q^1, a^1), \dots, (q^T, a^T) \rangle$. Fig. 1 (left side) shows a typical C that CONVEX handles, with $T = 5$ (six turns). Usually, q^0 is well-formed, and all other q^t are ad hoc.

Question. Each q^t is a sequence of words q_i^t , such that $q^t = \langle q_1^t, \dots, q_{|q^t|}^t \rangle$, where $|q^t|$ is the number of words in q^t . During answering, each word in q^t potentially maps to one or more items in K ($q_i^t \mapsto E \cup P \cup C \cup L$). However, since conversations revolve around entities of interest, we fixate on the mapped entities, and refer to them as $E(q^t)$. E.g., “Alan” in $q^1 \mapsto \{\text{Alan Arkin}\}$, and “score” in $q^2 \mapsto \{\text{composer, soundtrack, The Last Unicorn Soundtrack}\}$; so $E(q^1) = \{\text{Alan Arkin}\}$, and $E(q^2) = \{\text{The Last Unicorn Soundtrack}\}$.

Answer. Each answer a^t to question q^t is a (possibly multiple, single, or null-valued) set of entities or literals in K , i.e. $a^t \in E \cup L$ (questions asking for predicates or classes are usually not realistic). Each a^t is shaded in light blue in Fig. 1.

Context subgraph. In the CONVEX model, every turn t in C is associated with a context X^t , that is a subgraph grounded or anchored in a localized zone in K . Each X^t subgraph consists of: (i) the previous question entities in C , $\{E(q^1) \cup \dots \cup E(q^{t-1})\}$, (ii) previous answer entities in C : $\{a^1 \cup \dots \cup a^{t-1}\}$, (iii) intermediate nodes and edges connecting the above in K . All E nodes corresponding to turns $1, \dots, (t-1)$, are shaded in light green.

Frontier nodes. At every turn t , nodes in the k -hop neighborhood of X^t , $\mathcal{N}(X^t)$, define something like a border to which we may need to expand X^t for answering the next q^t (current nodes in X^t are subsumed in $\mathcal{N}(X^t)$). The number of hops k is small in practice, owing to the fact that typical users do not suddenly make large topic jumps during a specific C . Even then, since expanding X^t to include every $n \in \mathcal{N}(X^t)$ results in an exponential growth

rate for its size that we wish to avoid, we first select the best (top- r) nodes in $\mathcal{N}(X^t)$. These optimal expansion points in $\mathcal{N}(X^t)$ are referred to as *frontier nodes*, a ranked set $F^t = \{F_1^t, \dots, F_r^t\}$, and are the most relevant nodes with respect to the current question q^t and the current context X^t , as ranked by some *frontier score* (defined later). This entails that only those triples (along with qualifiers) (analogously, the resultant nodes and edges) that connect these F_i^t to the X^t are added to the context. The top-1 frontier node F_1^t at every t is shown in orange in the figure (multiple in case of ties).

Expanded context. Once all triples in K corresponding to frontier nodes F^t are added to X^t , we obtain an expanded context graph X_+^t . All nodes in X_+^t are candidate answers a_t , that are scored appropriately. Fig. 1 shows expanded contexts $\{X_+^t\}$ for every t in our example conversation. Corresponding X^t can be visualized by removing facts with the orange frontiers. Notably, $X_+^t = X^{t+1}$.

3 THE CONVEX ALGORITHM

We now describe the CONVEX conversation handler method, that can be envisaged as a seamless plug-in enabling stand-alone KG-QA systems to answer incomplete follow up questions with possibly ungrammatical and informal formulations. CONVEX thus requires an underlying KG, a standard QA system that can answer well-formulated questions, and the conversational utterances as input. On receiving an input question at a given turn, our method proceeds in two stages: (i) expand the context graph, and (ii) rank the answer candidates in the expanded graph. We discuss these steps next.

3.1 Context expansion

The initial question q^0 is answered by the KG-QA system that CONVEX augments, and say, that it produces answer(s) a^0 . Since entities in the original question are of prime importance in a conversation, we use any off-the-shelf named entity recognition and disambiguation (NERD) system like TagMe [11] or AIDA [14] to identify entities $E(q^0)$. Such $E(q^0)$, a^0 , and the KG connections between them initialize the context subgraph X^1 .

Now, when the first question q^1 arrives, we need to look for answer(s) a^1 in the vicinity of X^1 . The main premise of this work is not to treat *every node* in such neighborhood of X^1 , and more generally, X^t , as an answer candidate. This is because, over turns, expanding the context, by any means, is inevitable: users can freely drift away and revisit the initial entities of interest over the full course of the conversation. Under this postulate, the total number of such context nodes can easily go to the order of *millions*, aggravated by the presence of popular entities, especially countries (UK, Russia) or cities (Munich, Barcelona) in the KG around prominent entities of discussion (Harry Potter, Cristiano Ronaldo).

The logical course of action, then, is to perform this expansion in a somewhat austere fashion, which we propose to do as follows. We wish to identify some key nodes in the k -hop neighborhood of X^1 , that will prove the most worthwhile if included into X^1 (along with their connections to X^1) w.r.t. answering q^1 . We call these optimal expansion points *frontier nodes*. From here on, we outline frontier identification at a general conversation turn t , where $t = 0, 1, \dots, T$. Frontiers are marked by three signals: (i) relevance to the words in q^t ; (ii) relevance to the current context X^t ; and (iii) KG priors. We now explain these individual factors.

Relevance to question. The question words $\{q_i^t\}$ provide a direct clue to the relevant nodes in the neighborhood. However, there is often a vocabulary mismatch between what users specify in their questions and the KG terminology, as typical users are unaware of the KG schema. For example, let us consider $q^3 = \textit{Who did the score?}$. This indicates the sought information is about the score of the movie but unfortunately the KG does not use this term. So, we define the *matching similarity score* of a neighbor n with a question word using cosine similarity between word2vec [23] embeddings of the node label and the word. *Stopwords* like *and, of, to, etc.* are excluded from this similarity. For multiword phrases, we use an averaging of the word vectors [37]. The cosine similarity is originally in $[-1, +1]$: it is scaled to $[0, 1]$ using min-max normalization for comparability to the later measures. So we have:

$$\text{match}(n, q_i^t | n \in \mathcal{N}(X^t)) = \text{cos}_{\text{norm}}(\text{w2v}(\text{label}(n)), \text{w2v}(q_i^t)) \quad (1)$$

We then take the maximum of these word-wise scores to define the matching score of a candidate frontier to the question as a whole:

$$\text{match}(n, q^t | n \in \mathcal{N}(X^t)) = \max_i \text{match}(n, q_i^t) \quad (2)$$

Relevance to context. Nevertheless, such soft lexical matching with embeddings is hardly enough. Let us now consider the word “genre” in $q^3 = \textit{Genre of this band’s music?}$. Looking at the toy example in Fig. 2, we see that even with an exact match, there are five genre-s lurking in the vicinity at X^4 (there are several more in reality), where the one connected to *America* is the intended fit.

We thus define the relevance of a node n to the current context X^t as the total graph distance d_K (in number of hops in K) of n to the nodes in X^t . Note that we are interested in the relevance score being directly proportional to this quantity, and hence consider proximity, the reciprocal of distance ($1/d_K$), as the measure instead. For the aggregation over nodes x in X^t , we prefer $\sum(1/d_K(n, x))$ over $1/\sum(d_K(n, x))$ as the latter is more sensitive to outliers.

Next, not all nodes in X^t are valuable for the answering process. For anchoring a conversation C in a KG, entities that have appeared in *questions* or as *answers* in turns $0, \dots, (t-1)$, are what specifically




Figure 2: An illustration of the ambiguity in frontier node selection for a specific question word (“genre” in $q^4 = \textit{Genre of this band’s music?}$), and how effective scoring can potentially pick the best candidate in a noisy context graph X^4 .

matter. Thus, it suffices to consider only $\bigcup_{j=0}^{t-1} E(q^j)$ and $\bigcup_{j=0}^{t-1} a^j$ for computing the above proximities. We encode this factor using an indicator function $\mathbb{1}_{Q_A}(x|x \in X^t)$ that equals 1 if $x \in \bigcup_{j=0}^{t-1} [E(q^j) \cup a^j]$, and zero otherwise.

Contributions of such Q/A nodes $x \in X^t$ ($E(q^j)$ or a^j) should be weighted according to the turn in which they appeared in their respective roles, denoted by $\text{turn}(x)$. This is when such nodes had the “spotlight”, in a sense; so recent turns have higher weights than older ones. In addition, since the *entity in the first question* $E(q^0)$ may always be important as the theme of the conversation (The Last Unicorn), $\text{turn}(E(q^0))$ is set to the maximum value $(t-1)$ instead of zero. We thus define the *context proximity score* for neighbor n , normalized by the number of Q/A nodes in the context, as:

$$\text{prox}(n, X^t | n \in \mathcal{N}(X^t)) = \frac{\sum_{x \in X^t} \text{turn}(x) \cdot \mathbb{1}_{Q_A}(x) \cdot 1/d_K(n, x)}{\sum_{x \in X^t} \mathbb{1}_{Q_A}(x)} \quad (3)$$

KG priors. Finally, KG nodes have inherent salience (or prominence) values that reflect their likelihoods of being queried about in users’ questions. For example, Harry Potter has higher salience as opposed to some obscure book like Harry Black, and the same can be said to hold about the author predicate compared to has edition for books. Ideally, such salience should be quantified using large-scale query logs from real users that commercial Web search engines possess. In absence of such resources, we use a more intrinsic proxy for salience: the frequency of the concept in the KG. The raw frequency is normalized by corresponding maximum values for entities, predicates, classes, and literals, to give $\text{freq}_{\text{norm}}$. Thus, we have the *KG prior* for a node n as:

$$\text{prior}(n, K) = \text{freq}_{\text{norm}}(n, K) \quad (4)$$

Aggregation using Fagin’s Threshold Algorithm. We now have three independent signals from the question, context, and the KG regarding the likelihood of a node being a frontier at a given turn. We use the *Fagin’s Threshold Algorithm (FTA)* [10] to aggregate items in the three sorted lists $\mathcal{L} = \{\mathcal{L}_{\text{match}}, \mathcal{L}_{\text{prox}}, \mathcal{L}_{\text{prior}}\}$, that are created when candidates are scored by each of these signals. FTA is chosen as it is an optimal algorithm with correctness and performance guarantees for rank aggregation. In FTA, we perform

Algorithm 1: CONVEX ($K, T, q^0, a^0, E(q^0), \langle q^1, \dots, q^T \rangle, r$)

```

 $t = 1;$ 
initialize  $X^1 = E(q^0) \cup a^0 \cup K(E(q^0), a^0);$ 
while  $t \leq T$  do
  for  $n \in \mathcal{N}(X^t)$  do
    compute  $match(n, q^t)$  [Eq. 2];
    compute  $prox(n, X^t)$  [Eq. 3];
    compute  $prior(n, K)$  [Eq. 4];
    insert scores into sorted lists
       $\mathcal{L} = \{\mathcal{L}_{match}, \mathcal{L}_{prox}, \mathcal{L}_{prior}\};$ 
  end
  find  $\{F_i^t\}_{i=1}^r = \text{Fagin's-Threshold-Algorithm}(\mathcal{L}, r)$  [Eq. 5];
  assign  $E(q^t) = E(F^t);$ 
  expand  $X_+^t = X^t \cup facts(F^t);$ 
  for  $a \in X_+^t$  do
    compute  $ans-score(a)$  [Eq. 6];
  end
  find  $a^t = \arg \max_{X_+^t} ans-score(a);$ 
   $X^{t+1} \leftarrow X_+^t;$ 
   $t \leftarrow t + 1;$ 
end
Result:  $\langle a^1, \dots, a^T \rangle$ 

```

sorted access in parallel to each of the three sorted lists \mathcal{L}_i . As each candidate frontier node n is seen under sorted access, we retrieve each of its individual scores by random access. We then compute a frontier score for n as a simple linear combination of the outputs from Eqs. 2, 3, and 4 using frontier hyperparameters h_1^F , h_2^F , and h_3^F , where $\sum_i h_i^F = 1$:

$$\begin{aligned} & frontier(n, q^t | n \in \mathcal{N}(X^t)) \\ &= h_1^F \cdot match(n, q^t) + h_2^F \cdot prox(n, X^t) + h_3^F \cdot prior(n, K) \quad (5) \end{aligned}$$

In general, FTA requires a monotonic score aggregation function $frontier()$ such that $frontier(s_1, s_2, s_3) \leq frontier(s'_1, s'_2, s'_3)$ whenever $s_i \leq s'_i, \forall i$, where the component scores of $frontier(n)$ are denoted as s_i 's (Eq. 5) in corresponding lists \mathcal{L}_i . Once the above is done, as nodes are accessed from \mathcal{L}_i , if this is one of the top r answers so far, we remember it. Here, we assume a buffer of bounded size. For each \mathcal{L}_i , let \hat{s}_i be the score of the last node seen under sorted access. We define the threshold value $thresh$ to be $frontier(\hat{s}_1, \hat{s}_2, \hat{s}_3)$. When r nodes have been seen whose frontier score is at least $thresh$, then we stop and return the top r nodes as the final frontiers.

Thus, at the end of this step, we have a set of r frontier nodes $\{F_i^t\}_{i=1}^r$ for turn t . If any of these frontiers are entities, they are used to populate $E(q^t)$. We add the triples (along with qualifiers if any) that connect the F_i^t to the current context X^t to produce the expanded graph X_+^t (the step containing $facts(F^t)$ in Algorithm 1).

3.2 Answer ranking

Our task now is to look for answers a_t in X_+^t . Since frontiers are the most relevant nodes in X_+^t w.r.t question q_t , it is expected that a_t will appear in their close proximity.

Attribute	Value
Title	Generate question-answer conversations on popular entities
Description	Generate conversations in different domains (books, movies, music, soccer, and TV series) on popular entities of your choice. You need to ask natural questions and provide the corresponding answers via Web search.
Total participants	70
Time allotted per HIT	6 hours
Time taken per HIT	3 hours
Payment per HIT	25 Euros

Table 2: Basic details of the AMT HIT (five conversations).

However, labels of frontier nodes only reflect what was *explicit* in q^t , the unspecified or *implicit* part of the context in q^t usually refers to a previous question or answer entity ($\bigcup_{j=0}^{t-1} [E(q^j) \cup a^j]$). Thus, we should consider closeness to these context entities as well. Note that just as before, frontiers and Q/A entities both come with corresponding weights: frontier scores and turn id's, respectively. Thus, while considering proximities is key here, using weighted versions is a more informed choice. We thus score every node $a \in X_+^t$ by its weighted proximity, using Eqs. 3 and 5, as follows (again, we invert distance to use a measure *directly proportional* to the candidacy of an answer node):

$$\begin{aligned} a^t = \arg \max_{a \in X_+^t} & \left[h_1^A \cdot \frac{\sum_{i=1}^r [frontier(F_i^t) \cdot (1/d_K(a, F_i^t))]}{r} \right. \\ & \left. + h_2^A \cdot \frac{\sum_{x \in X^t} turn(x) \cdot \mathbb{1}_{QA}(x) \cdot 1/d_K(a, x)}{\sum_{x \in X^t} \mathbb{1}_{QA}(x)} \right] \quad (6) \end{aligned}$$

Contributions by proximities to frontier and Q/A nodes (each normalized appropriately) are again combined linearly with answer hyperparameters h_1^A and h_2^A , where $h_1^A + h_2^A = 1$. Thus, the final answer score also lies in $[0, 1]$. Finally, the top scoring a^t (possibly multiple, in case of ties) node(s) is returned as the *answer* to q^t .

The CONVEX method is outlined in Algorithm 1. As mentioned before, a^0 and $E(q^0)$ are obtained by passing q^0 through a stand-alone KG-QA system, and a NERD algorithm, respectively. $K(\cdot)$ returns all KG triples that contain the arguments of this function, and the generalized $E(\cdot)$ returns the set of entities from its arguments. Note that this algorithm illustrates the workings of CONVEX in a static setting when all $q^t (t = 0 \dots T)$ are given upfront; in a real setting, each q^t is issued interactively with a user in the loop.

4 THE CONVQUESTIONS BENCHMARK

4.1 Benchmark creation

Limitations of current choices. Popular benchmarks for KG-QA like WebQuestions [5], SimpleQuestions [6], WikiMovies [24], ComplexWebQuestions [34], and ComQA [3], are all designed for one-shot answering with well-formulated questions. The CSQA dataset [30] takes preliminary steps towards the sequential KG-QA paradigm, but it is extremely artificial: initial and follow-up questions are generated semi-automatically via templates, and sequential utterances are only simulated by stitching questions with shared entities or relations in a thread, without a logical flow. QBLink [9],

Turn	Books	Movies	Soccer	Music	TV series
q^0	When was the first book of the book series <i>The Dwarves</i> published?	Who played the joker in <i>The Dark Knight</i> ?	Which European team did Diego Costa represent in the year 2018?	Led Zeppelin had how many band members?	Who is the actor of James Gordon in <i>Gotham</i> ?
a^0	2003	Heath Ledger	Atlético Madrid	4	Ben McKenzie
q^1	What is the name of the second book?	When did he die?	Did they win the Super Cup the previous year?	Which was released first: <i>Houses of the Holy</i> or <i>Physical Graffiti</i> ?	What about <i>Bullock</i> ?
a^1	The War of the Dwarves	22 January 2008	No	Houses of the Holy	Donal Logue
q^2	Who is the author?	Batman actor?	Which club was the winner?	Is the rain song and immigrant song there?	Creator?
a^2	Markus Heitz	Christian Bale	Real Madrid C.F.	No	Bruno Heller
q^3	In which city was he born?	Director?	Which English club did Costa play for before returning to Atlético Madrid?	Who wrote those songs?	Married to in 2017?
a^3	Homburg	Christopher Nolan	Chelsea F.C.	Jimmy Page	Miranda Cowley
q^4	When was he born?	Sequel name?	Which stadium is this club's home ground?	Name of his previous band?	Wedding date first wife?
a^4	10 October 1971	The Dark Knight Rises	Stamford Bridge	The Yardbirds	19 June 1993

Table 3: Representative conversations in CONVQUESTIONS from each domain, highlighting the stiff challenges they pose.

CoQA [27], ans ShARC [29] are recent resources for sequential QA over text. The SQA resource [16], derived from WikiTableQuestions [25], is aimed at driving conversational QA over (relatively small) Web tables.

Conceptual challenges. In light of such limitations, we overcome several conceptual challenges to build the first realistic benchmark for conversational KG-QA, anchored in Wikidata. The key questions included, among others: Should we choose q^0 from existing benchmarks and ask humans to create only follow-ups? Should the answers already come from some KG-QA system, observing which, users create follow-ups? Should we allocate templates to crowdworkers to systematically generate questions that miss either entities, predicates, and types? Can we interleave questions by different workers to create a large number of conversations? Can we permute the order of follow-ups to generate an even larger volume? If there are multiple correct a^t , and q^{t+1} in the benchmark involves a different a^t than what the system returns at run-time, how can we evaluate such a dynamic workflow? How can we build a KG-QA resource that is faithful to the setup but is not overly limited to the information the KG contains today?

Creating CONVQUESTIONS. With insights from a meticulous in-house pilot study with ten students over two weeks, we posed the conversation generation task on Amazon Mechanical Turk (AMT) in the most *natural setup*: Each crowdworker was asked to build a conversation by asking five sequential questions *starting* from any *seed* entity of his/her choice, as this is an intuitive mental model that humans may have when satisfying their real information needs via their search assistants. A system-in-the-loop is hardly ideal: this creates comparison across methods challenging, is limited by the shortcomings of the chosen system, and most crucially, there exist no such systems today with satisfactory performance. In a single AMT Human Intelligence Task (HIT), Turkers had to create one conversation each from five domains: “Books”, “Movies”, “Soccer”, “Music”, and “TV Series” (other potential choices were “Politics”, but we found that it quickly becomes subjective, and “Finance”,

but that is best handled by relational databases and not curated KGs). Each conversation was to have five turns, including q^0 . To keep conversations as *natural* as possible, we neither interleaved questions from multiple Turkers, nor permuted orders of questions within a conversation. For quality, only AMT Master Workers (who have consistently high performances: see https://www.mturk.com/help#what_are_masters), were allowed to participate. We registered 70 participants, and this resulted in 350 initial conversations, 70 from each domain.

Along with questions, Turkers were asked to provide textual surface forms and Wikidata links of the seed entities and the *answers* (via Web search), along with *paraphrases* of each question. The paraphrases provided us with two versions of the same question, and hence a means of augmenting the core data with several interesting variations that can simultaneously boost and test the robustness of KG-QA systems [8]. Since paraphrases of questions (any q^t) are always semantically equivalent and interchangeable, each conversation with five turns thus resulted in $2^5 = 32$ distinct conversations (note that this *does not* entail shuffling sequences of the utterances). Thereby, in total, we obtained $350 \times 32 = 11,200$ such conversations, that we release with this paper.

If the answers were dates or literals like measurable quantities with units, Turkers were asked to follow the Wikidata formats for the same. They were provided with *minimal syntactic guidelines* to remain natural in their questions. They were shown judiciously selected examples so as not to ask opinionated questions (like *best film by this actor?*), or other non-factoid questions (causal, procedural, etc.). The authors invested substantial manual effort for quality control and spam prevention, by verifying both answers of random utterances, and alignments between provided texts and Wikidata URLs. Each question was allowed to have at most three answers, but single-answer questions were encouraged to preclude the possibility of non-deterministic workflows during evaluation.

To allow for CONVQUESTIONS being relevant for a few years into the future, we encouraged users to ask complex questions involving

joins, comparisons, aggregations, temporal information needs, and so on. Given the complexity arising from incomplete cues, these additional facets pose an even greater challenge for future KG-QA systems. So as not to restrict questions to only those predicates that are present in Wikidata today, relations connecting question and answer entities are sometimes missing in the KG but can be located in sources like Wikipedia, allowing scope for both future growth of the KG, and experimentation with text plus KG combinations.

4.2 Benchmark analysis

Basic details of our AMT HIT are provided in Table 2 for reference. Question entities and expected answers had a balanced distribution among human (actors, authors, artists) and non-human types (books, movies, stadiums). Detailed distributions are omitted due to lack of space. Illustrative examples of challenging questions from CONVQUESTIONS are in Table 3. We see manifestations of: incomplete cues (TV Series; q^3), ordinal questions (Books; q^1), comparatives (Music; q^1), indirections (Soccer; q^4), anaphora (Music; q^3), existentials (Soccer; q^2), temporal reasoning (Soccer; q^3), among other challenges. The average lengths of the first and follow-up questions were 9.07 and 6.20 words, respectively. Finally, we present the key quantifier for the difficulty in our benchmark: the average KG distance of answers from the original seed entity is 2.30 hops, while the highest goes up to as high as *five* KG hops. Thus, an approach that remains fixated on a specific entity is doomed to fail: context expansion is the key to success on CONVQUESTIONS.

5 EXPERIMENTAL SETUP

5.1 Baselines and Metrics

Stand-alone systems. We use the state-of-the-art system QAnswer [7], and also Platypus [35], as our stand-alone KG-QA systems, that serve as baselines, and which we enhance with CONVEX. At the time of writing (May 2019), these are the only two systems that have running prototypes over Wikidata.

To make CONVEX a self-sufficient system, we also implement a *naïve* version of answering the first question as follows. Entities are detected in q^0 using the TagMe NERD system [11], and mapped to their Wikidata IDs via Wikipedia links provided by TagMe. Embeddings were obtained by averaging word2vec vectors of the non-entity words in q^0 , and their cosine similarities were computed around each of the predicates around the detected entities $E(q^0)$. Finally, the best (E, P) pair was found (as a joint disambiguation), and the returned answer was the subject or object in the triple according as the triple structure was $\langle \cdot, P, E \rangle$ or $\langle E, P, \cdot \rangle$.

Due to the complexity in even the first question q^0 in CONVQUESTIONS, all of the above systems achieve a very poor performance for a^0 on the benchmark. This limits the value that CONVEX can help these systems achieve, as $E(q^0)$ and a^0 together initialize X^1 . To decouple the effect of the original QA system, we experiment with an *Oracle* strategy, where we use $E(q^0)$ and a^0 provided by the human annotator (Turker who created the conversation).

Conversation models. As intuitive alternative strategies to CONVEX for handling conversations, we explore two variants: (i) the *star-join*, and (ii) the *chain-join* models. The naming is inspired by DB terminology, where a star query looks for a join on several attributes around a single variable (of the form `SELECT ?x WHERE`

`{?x att1 val1 . ?x att2 val2 . ?x att3 val3}}`), while a chain SQL searches for a multi-variable join via indirections (`SELECT ?x WHERE {?x att1 ?y . ?y att2 ?z . ?z att3 val1}}`). For conversations, this entails the following: in the star model, the entity in q^0 is always assumed to be the entity in all subsequent utterances (like The Last Unicorn). The best predicate is disambiguated via a search around such $E(q^0)$ using similarities of word2vec embeddings of Wikidata phrases and non-entity words in q^0 . The corresponding missing argument from the triple is returned as the answer. In the chain model of a conversation, the previous answer a^{t-1} is always taken as the reference entity at turn t , instead of $E(q^0)$. Predicate selection and answer detection are done analogously as in the star model.

No frontiers. We also investigated whether the idea of a frontier node in itself was necessary, by defining an alternative configuration where we *optimize an answer-scoring objective directly*. The same three signals of question matching, context proximity, and KG priors were aggregated (Eqs. 2, 3, and 4), and the Fagin’s Threshold Algorithm was again applied for obtaining the top- r list. However, these top- r returned nodes are now directly the answers. The process used translates to a *branch-and-bound strategy* for iteratively exploring the neighborhood of the initial context ($E(q^0)$, a^0 , and their interconnections) as follows, without explicitly materializing a context subgraph. The 2-hop neighborhood (2-hop as we now directly score for an answer, without finding a frontier first) of each node in the context at a given turn is scored on its likelihood of being an answer, in a breadth-first manner. The first computed score defines a lower bound on the node being a potential answer, that is updated as better candidates are found. If a node’s answer score is lower than the lower bound so far, it is not expanded further (its neighborhood is not explored anymore). We keep exploring the 2-hop neighborhood of the context iteratively until we do not find any node in K better than the current best answer.

End-to-end neural model. We compared our results with D2A (Dialog-to-Action) [12], the state-of-the-art end-to-end neural model for conversational KG-QA. Since CONVEX is an unsupervised method that does not rely on training data, we used the D2A model pre-trained on the large CSQA benchmark [30]. D2A manages dialogue memory using a generative model based on a flexible grammar.

Question completion. An interesting question to ask at this stage is whether an attempt towards completing the follow-up utterances is worthwhile. While a direct adaptation of a method like [17] is infeasible due to absence of training pairs and the need for *rewriting* as opposed to plain *completion*, we investigate certain reasonable alternatives: (i) when q^t is concatenated with keywords (all nouns and verbs) from q^0 ; (ii) when q^t is concatenated with a^0 ; (iii) when q^t is concatenated with keywords from q^{t-1} ; and, (iv) with a^{i-1} . These variants are then passed through the stand-alone KG-QA system. Fortunately, the state-of-the-art system QAnswer is totally syntax-agnostic, and searches the KG with all question cue words to formulate an optimal SPARQL query whose components best *cover* the mapped KG items. This syntax-independent approach was vital as it would be futile to massage the “completed” questions above into grammatically correct forms. Platypus, on the other hand, is totally dependent on an accurate dependency parse of the input utterance, and hence is unsuitable for plugging in these question completion strategies.

Metrics. Since most questions in CONVQUESTIONS had exactly one or at most a few correct answers, we used the standard metrics of Precision at the top rank (P@1), Mean Reciprocal Rank (MRR), and Hit@5 metrics. The last measures the fraction of times a correct answer was retrieved within the top-5 positions.

5.2 Configuration

Dataset. We evaluate CONVEX and other baselines on CONVQUESTIONS. A random 20% of the 11k conversations was held out for tuning model parameters, and the remaining 80% was used for testing. Care was taken that this development set was generated from a separate set of seed conversations (70 out of the original 350) so as to preclude possibilities of “leakage” on to the test set.

Initialization. We use Wikidata (www.wikidata.org) as our underlying KG, and use the complete RDF dump in NTriples format from 15 April 2019 (<http://bit.ly/2QhsSDC>, ≈ 1.3 TB uncompressed). Identifier triples like those containing predicates like Freebase ID, IMDb ID, etc. were excluded. We used indexing with HDT (www.rdfhdt.org/) that enables much faster lookups. The Python library NetworkX (<https://networkx.github.io/>) was used for graph processing. TagMe was used for NERD, and word2vec embeddings were obtained via the *spaCy* package. Stanford CoreNLP [22] was used for POS tagging to extract nouns and verbs for question completion. The ideal number of frontier nodes, r , was found to be **three** by tuning on the dev set.

6 RESULTS AND INSIGHTS

6.1 Key findings

Table 4 lists main results, where all configurations *are run on the follow-up utterances* in the CONVQUESTIONS test (8,960 conversations; 35,840 questions). An asterisk (*) indicates statistical significance of CONVEX-enabled systems over the strongest baseline in the group, under the 2-tailed paired t -test at $p < 0.05$ level. We make the following key observations.

CONVEX enables stand-alone systems. The state-of-the-art QAnswer [7] scores only about 0.011 – 0.064 (since it produces sets and not ranked lists, all metric values are identical) on its own on the incomplete utterances, which it is clearly not capable of addressing. When CONVEX is applied, its performance jumps significantly to 0.172 – 0.264 (MRR) across the domains. We have exactly the same trends with the Platypus system. The naive strategy with direct entity and predicate linking performs hopelessly in absence of explicit cues, but with CONVEX we again see noticeable improvements, brought in by a relevant context graph and its iterative expansion. In the Oracle method, a^0 is known, and hence a row by itself is not meaningful. However, contrasting Oracle+CONVEX with other “+CONVEX” methods, we see that there is significant room for improvement that can be achieved by answering q^0 correctly.

Star and chain models of conversations fall short. For every configuration, we see the across-the-board superiority of CONVEX-boosted methods over star- and chain-models (often over 100% gains). This clearly indicates that while these are intuitive ways of modeling human conversation (as seen in the often respectable values that these achieve), they are insufficient and oversimplified. Evidently, real humans rather prefer the middle path: sometimes hovering around the initial entity, sometimes drifting in a chain of

answers. A core component of CONVEX that we can attribute this pattern to, is the turn-based weighting of answer and context proximity that prefers entities in the first and the last turns. “QAnswer + Star” and “Platypus + Star” achieve the same values as they both operate around the same q^0 entity detected by TagMe.

CONVEX generalizes across domains. In Table 4, we also note that the performance of CONVEX stretches across all five domains (even though the nature of questions in each of these domains have their own peculiarities), showing the potential of our unsupervised approach in new domains with little training resources, or to deal with cold starts in enterprise applications. While we did tune hyperparameters individually for each domain, there were surprisingly little variation across them ($h_1^F \approx 0.5 - 0.6$, $h_2^F \approx 0.3 - 0.4$, $h_3^F \approx 0.1$, $h_1^A \approx 0.8 - 0.9$, $h_2^A \approx 0.1 - 0.2$).

Frontiers help. We applied our frontier-less approach over the oracle annotations for q^0 , and in the row marked “Oracle + No frontiers” in Table 4, we find that this results in degraded performance. We thus claim that locating frontiers is an essential step before answer detection. The primary reason behind this is that answers only have low direct matching similarity to the question, making a 2-stage approach worthwhile. Also, exploring a 2-hop neighborhood was generally found to suffice: nodes further away from the initial context rarely manage to “win”, due to the proximity score component quickly falling off as KG-hops increase.

Pre-trained models do not suffice. D2A produces a single answer for every utterance, which is why the three metrics are equal. From the D2A row in Table 4, we observe that pre-trained neural models do not work well off-the-shelf on CONVQUESTIONS (when compared to the CONVEX-enabled QAnswer row, for example). This is mostly due to the restrictive patterns in the CSQA dataset, owing to its semi-synthetic mode of creation. A direct comparison, though, is not fair, as CONVEX is an enabler method for a stand-alone KG-QA system, while D2A is an end-to-end model. Nevertheless, the main classes of errors come from: (i) a predicate necessary in CONVQUESTIONS that is absent in CSQA (D2A cannot answer temporal questions like *In what year was Ender’s game written?* as such relations are absent in CSQA); (ii) D2A cannot generate 2-hop KG triple patterns; (iii) D2A cannot resolve long-term co-references in questions (pronouns only come from the last turn in CSQA, but not in CONVQUESTIONS); (iv) In CSQA, co-references are almost always indicated as “it” or “that one”. But since CONVQUESTIONS is completely user-generated, we have more challenging cases with “this book”, “the author”, “that year”, and so on.

CONVEX outperforms question completion methods. Comparison with question completion methods are presented in Table 5. Clear trends show that while these strategies generally perform better than stand-alone systems (contrasting QAnswer with Table 4, for say, Movies, we see 0.050 – 0.109 vs. 0.032 on MRR previously), use of CONVEX results in higher improvement (0.264 MRR on Movies). This implies that question completion is hardly worthwhile in this setup when the KG structure already reveals a great deal about the underlying user intents left implicit in follow-up utterances.

6.2 Analysis

CONVEX maintains its performance over turns. One of the most promising results of this zoomed-in analysis is that the MRR

Domain	Movies			TV Series			Music			Books			Soccer		
Method	P@1	MRR	Hit@5	P@1	MRR	Hit@5	P@1	MRR	Hit@5	P@1	MRR	Hit@5	P@1	MRR	Hit@5
QAnswer [7]	0.032	0.032	0.032	0.064	0.064	0.064	0.020	0.020	0.020	0.011	0.011	0.011	0.020	0.020	0.020
QAnswer + CONVEX	0.222*	0.264*	0.311*	0.136*	0.172*	0.214*	0.168	0.197*	0.232*	0.177	0.213*	0.252*	0.179*	0.221*	0.265*
QAnswer + Star	0.201	0.201	0.201	0.132	0.132	0.132	0.183	0.183	0.183	0.199	0.199	0.199	0.170	0.170	0.170
QAnswer + Chain	0.077	0.077	0.077	0.028	0.028	0.028	0.056	0.056	0.056	0.034	0.034	0.034	0.044	0.044	0.044
Platypus [35]	0.000	0.000	0.000	0.000	0.000	0.000	0.005	0.005	0.005	0.002	0.002	0.002	0.004	0.004	0.004
Platypus + CONVEX	0.218*	0.255*	0.295*	0.124	0.153*	0.189*	0.167	0.197*	0.233*	0.180	0.216*	0.256*	0.179*	0.222*	0.269*
Platypus + Star	0.201	0.201	0.201	0.132	0.132	0.132	0.183	0.183	0.183	0.199	0.199	0.199	0.170	0.170	0.170
Platypus + Chain	0.047	0.047	0.047	0.000	0.000	0.000	0.028	0.028	0.028	0.028	0.028	0.028	0.015	0.015	0.015
Naive	0.016	0.016	0.016	0.020	0.020	0.020	0.021	0.021	0.021	0.007	0.007	0.007	0.016	0.016	0.016
Naive + CONVEX	0.212*	0.252*	0.296*	0.121	0.149*	0.185*	0.164	0.194*	0.229*	0.176	0.210*	0.248*	0.161*	0.201*	0.245*
Naive + Star	0.205	0.205	0.205	0.129	0.129	0.129	0.185	0.185	0.185	0.205	0.205	0.205	0.154	0.154	0.154
Naive + Chain	0.059	0.059	0.059	0.014	0.014	0.014	0.039	0.039	0.039	0.051	0.051	0.051	0.031	0.031	0.031
Oracle + CONVEX	0.259*	0.305*	0.355*	0.178	0.218*	0.269*	0.190	0.237	0.293*	0.198	0.246*	0.303*	0.188*	0.234*	0.284*
Oracle + Star	0.257	0.257	0.257	0.194	0.194	0.194	0.241	0.241	0.241	0.241	0.241	0.241	0.179	0.179	0.179
Oracle + Chain	0.094	0.094	0.094	0.031	0.031	0.031	0.040	0.040	0.040	0.053	0.053	0.053	0.016	0.016	0.016
Oracle + No frontiers	0.124	0.153	0.191	0.073	0.094	0.125	0.116	0.144	0.185	0.103	0.137	0.199	0.087	0.122	0.166
D2A [12]	0.090	0.090	0.090	0.067	0.067	0.067	0.072	0.072	0.072	0.121	0.121	0.121	0.107	0.107	0.107

The highest value in a group (metric-domain-system triple) is in **bold**. QAnswer and Platypus return only a top-1 answer and not ranked lists, and hence have the same P@1, MRR, and Hit@5 values.

Table 4: Our main results on follow-up utterances in CONVQUESTIONS showing how CONVEX enables KG-QA enables for conversations, and its comparison with baselines.

Method	Movies	TV	Music	Books	Soccer
QAnswer + CONVEX	0.264*	0.172*	0.197*	0.213*	0.221*
QAnswer + q^0 keywords	0.071	0.052	0.084	0.039	0.075
QAnswer + a^0	0.077	0.054	0.048	0.096	0.045
QAnswer + q^{t-1} keywords	0.050	0.045	0.045	0.025	0.046
QAnswer + a^{t-1}	0.109	0.079	0.093	0.064	0.070

Table 5: Comparison with question completion strategies (MRR). The highest value in a column is in bold.

Turn	Movies	TV	Music	Books	Soccer
1	0.375	0.393	0.080	0.446	0.357
2	0.375	0.250	0.214	0.281	0.188
3	0.161	0.205	0.124	0.435	0.304
4	0.325	0.214	0.044	0.375	0.137

Table 6: Performance of CONVEX over turns (MRR).

for CONVEX (measured via its combination with the Oracle, to decouple the effect of the QA system) does not diminish over turns. This shows particular robustness of our graph-based method: while we may produce several wrong results during the session of the conversation, we are not bogged down by any single mistake, as the context graph retains several scored candidates within itself, guarding against “near misses”. This is in stark contrast to the chain model, where it is exclusively dependent on a^{t-1} .

Error analysis. CONVEX has two main steps in its pipeline: context expansion, and answer ranking. Analogously, there are two main cases of error: when the answer is not pulled in when X^t is expanded at the frontiers (incorrect frontier scoring), or when the answer is there in X_+^t but is not ranked at the top. These numbers are shown in Table 7. We find that there is significant scope for improvement for frontier expansion, as 80 – 90% errors lie in this bag. It is however, heartening to see that no particular turn is singularly affected. This calls for more informed frontier scoring than our current strategy. Answer ranking can be improved with

Scenario	Turn 1	Turn 2	Turn 3	Turn 4
Ans. not in expanded graph	87.1	79.8	89.2	89.6
Ans. in expanded graph but not in top-1	12.9	20.2	10.8	10.4

Table 7: Error analysis (percentages of total errors).

Utterance: <i>What was the name of the director?</i> (Movies, Turn 4)	
Intent: <i>Who was the director of the movie Breakfast at Tiffany’s?</i>	
Utterance: <i>What about Mr Morningstar?</i> (TV Series, Turn 2)	
Intent: <i>Which actor plays the role of Mr Morningstar in the TV series Lucifer?</i>	
Utterance: <i>What record label put out the album?</i> (Music, Turn 3)	
Intent: <i>What is the name of the record label of the album Cosmic Thing?</i>	
Utterance: <i>written in country?</i> (Books, Turn 4)	
Intent: <i>In which country was the book “The Body in the Library” by Agatha Christie written?</i>	
Utterance: <i>Who won the World Cup that year?</i> (Soccer, Turn 4)	
Intent: <i>Which national team won the 2010 FIFA World Cup?</i>	

Table 8: Representative examples where Oracle + CONVEX produced the best answer at the top-1, but neither Oracle + Star, nor Oracle + Chain could.

better ways of aggregating the two proximity signals. Table 8 lists anecdotal examples of success cases with CONVEX.

7 RELATED WORK

Question answering over KGs. Starting with early approaches in 2012-’13 [5, 36, 38], based on parsing questions via handcoded templates and grammars, KG-QA already has a rich body of literature. While templates continued to be a strong line of work due to its focus on interpretability and generalizability [1, 2, 4, 7, 35], a parallel thread has focused on neural methods driven by performance gains [15, 20, 31]. Newer trends include shifts towards more complex questions [19, 21, 34], and fusion of knowledge graphs and text [31, 33]. However, none of these approaches can deal with incomplete questions in a conversational setting.

Conversational question answering. Saha et al. [30] introduce the paradigm of sequential question answering over KGs, and create a large benchmark CSQA for the task, along with a baseline

with memory networks. Guo et al. [12] propose D2A, an end-to-end technique for conversational KG-QA, that introduces dialog memory management for inferring the logical form of current utterances. While our goal is rather to build a conversation enabler method, we still compare with, and outperform the CSQA-trained D2A model on CONVQUESTIONS.

Question completion approaches [17, 26, 28] target this setting by attempting to create full-fledged interrogatives from partial utterances while being independent of the answering resource, but suffer in situations without training pairs and with ad hoc styles. Nevertheless, we try to compare with this line of thought, and show that such completion may not be necessary if the underlying KG can be properly exploited.

Iyyer et al. [16] initiate the direction of sequential QA over tables using dynamic neural semantic parsing trained via weakly supervised reward-guided search, and evaluate by decomposing a previous benchmark of complex questions [25] to create sequential utterances. However, such table-cell search methods cannot scale to real-world, large-scale curated KGs.

QBLINK [9], CoQA [27], and SHARC [29] are recent benchmarks aimed at driving conversational QA over text, and the allied paradigm in text comprehension on interactive QA [18]. Hixon et al. [13] try to learn concept knowledge graphs from conversational dialogues over science questions, but such KGs are fundamentally different from curated ones like Wikidata with millions of facts.

8 CONCLUSION

Through CONVEX, we showed how judicious graph expansion strategies with informed look-ahead, can help stand-alone KG-QA systems cope with some of the challenges posed by incomplete and ad hoc follow-up questions in fact-centric conversations. CONVEX is completely unsupervised, and thus can be readily applied to new domains, or be deployed in enterprise setups with little or no training data. Further, being a graph-based method, each turn and the associated expansions can be easily visualized, resulting in interpretable evidence for answer derivation: an unsolved concern for many neural methods for QA. Nevertheless, CONVEX is just a first step towards solving the challenge of conversational KG-QA. We believe that the CONVQUESTIONS benchmark, reflecting real user behavior, can play a key role in driving further progress.

Acknowledgements. We sincerely thank Daya Guo (Sun Yat-sen University) for his help in executing D2A on CONVQUESTIONS.

REFERENCES

- [1] Abdalghani Abujabal, Rishiraj Saha Roy, Mohamed Yahya, and Gerhard Weikum. 2017. QUINT: Interpretable Question Answering over Knowledge Bases. In *EMNLP*.
- [2] Abdalghani Abujabal, Rishiraj Saha Roy, Mohamed Yahya, and Gerhard Weikum. 2018. Never-ending learning for open-domain question answering over knowledge bases. In *WWW*.
- [3] Abdalghani Abujabal, Rishiraj Saha Roy, Mohamed Yahya, and Gerhard Weikum. 2019. ComQA: A Community-sourced Dataset for Complex Factoid Question Answering with Paraphrase Clusters. In *NAACL-HLT '19*.
- [4] Hannah Bast and Elmar Haussmann. 2015. More accurate question answering on Freebase. In *CIKM*.
- [5] Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on Freebase from question-answer pairs. In *EMNLP*.
- [6] Antoine Bordes, Nicolas Usunier, Sumit Chopra, and Jason Weston. 2015. Large-scale simple question answering with memory networks. In *arXiv*.
- [7] Dennis Diefenbach, Pedro Henrique Migliatti, Omar Qawasme, Vincent Lully, Kamal Singh, and Pierre Maret. 2019. QAnswer: A Question Answering prototype

- bridging the gap between a considerable part of the LOD cloud and end-users. In *WWW*.
- [8] Li Dong, Jonathan Mallinson, Siva Reddy, and Mirella Lapata. 2017. Learning to Paraphrase for Question Answering. In *EMNLP*.
- [9] Ahmed Elgohary, Chen Zhao, and Jordan Boyd-Graber. 2018. A dataset and baselines for sequential open-domain question answering. In *EMNLP*.
- [10] Ronald Fagin, Amnon Lotem, and Moni Naor. 2003. Optimal aggregation algorithms for middleware. *Journal of Computer and System Sciences* 66, 4 (2003).
- [11] Paolo Ferragina and Ugo Scaiella. 2010. TAGME: On-the-fly annotation of short text fragments (by Wikipedia entities). In *CIKM*. 1625–1628.
- [12] Daya Guo, Duyu Tang, Nan Duan, Ming Zhou, and Jian Yin. 2018. Dialog-to-action: Conversational question answering over a large-scale knowledge base. In *NeurIPS*.
- [13] Ben Hixon, Peter Clark, and Hannaneh Hajishirzi. 2015. Learning knowledge graphs for question answering through conversational dialog. In *NAACL-HLT*.
- [14] Johannes Hoffart, Mohamed Amir Yosef, Ilaria Bordino, Hagen Fürstenau, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 2011. Robust disambiguation of named entities in text. In *EMNLP*.
- [15] Xiao Huang, Jingyuan Zhang, Dingcheng Li, and Ping Li. 2019. Knowledge graph embedding based question answering. In *WSDM*.
- [16] Mohit Iyyer, Wen-tau Yih, and Ming-Wei Chang. 2017. Search-based neural structured learning for sequential question answering. In *ACL*.
- [17] Vineet Kumar and Sachindra Joshi. 2017. Incomplete follow-up question resolution using retrieval based sequence to sequence learning. In *SIGIR*.
- [18] Huayu Li, Martin Renqiang Min, Yong Ge, and Asim Kadav. 2017. A context-aware attention network for interactive question answering. In *KDD*.
- [19] Xiaolu Lu, Soumajit Pramanik, Rishiraj Saha Roy, Abdalghani Abujabal, Yafang Wang, and Gerhard Weikum. 2019. Answering Complex Questions by Joining Multi-Document Evidence with Quasi Knowledge Graphs. In *SIGIR*.
- [20] Denis Lukovnikov, Asja Fischer, Jens Lehmann, and Sören Auer. 2017. Neural Network-based Question Answering over Knowledge Graphs on Word and Character Level. In *WWW*.
- [21] Kangqi Luo, Fengli Lin, Xusheng Luo, and Kenny Zhu. 2018. Knowledge Base Question Answering via Encoding of Complex Query Graphs. In *EMNLP*.
- [22] Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *ACL*.
- [23] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NeurIPS*.
- [24] Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. Key-value memory networks for directly reading documents. In *arXiv*.
- [25] Panupong Pasupat and Percy Liang. 2015. Compositional Semantic Parsing on Semi-Structured Tables. In *ACL*.
- [26] Dinesh Raghu, Sathish Indurthi, Jitendra Ajmera, and Sachindra Joshi. 2015. A statistical approach for non-sentential utterance resolution for interactive QA system. In *SIGDIAL*.
- [27] Siva Reddy, Danqi Chen, and Christopher D Manning. 2018. CoQA: A conversational question answering challenge.
- [28] Gary Ren, Xiaochuan Ni, Manish Malik, and Qifa Ke. 2018. Conversational Query Understanding Using Sequence to Sequence Modeling. In *WWW*.
- [29] Marzieh Saeidi, Max Bartolo, Patrick Lewis, Sameer Singh, Tim Rocktäschel, Mike Sheldou, Guillaume Bouchard, and Sebastian Riedel. 2018. Interpretation of Natural Language Rules in Conversational Machine Reading. In *EMNLP*.
- [30] Amrita Saha, Vardaan Pahuja, Mitesh M Khapra, Karthik Sankaranarayanan, and Sarath Chandar. 2018. Complex sequential question answering: Towards learning to converse over linked question answer pairs with a knowledge graph. In *AAAI*.
- [31] Uma Sawant, Saurabh Garg, Soumen Chakrabarti, and Ganesh Ramakrishnan. 2019. Neural architecture for question answering using a knowledge graph and web corpus. *Information Retrieval Journal* (2019).
- [32] Fabian Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. YAGO: A core of semantic knowledge. In *WWW*.
- [33] Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William Cohen. 2018. Open Domain Question Answering Using Early Fusion of Knowledge Bases and Text. In *EMNLP*.
- [34] Alon Talmor and Jonathan Berant. 2018. The Web as a Knowledge-Base for Answering Complex Questions. In *NAACL-HLT*.
- [35] Thomas Pellissier Tanon, Marcos Dias de Assuncao, Eddy Caron, and Fabian M Suchanek. 2018. Demoing Platypus – A multilingual question answering platform for Wikidata. In *ESWC*.
- [36] 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*.
- [37] John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2016. Towards universal paraphrastic sentence embeddings. In *ICLR*.
- [38] Mohamed Yahya, Klaus Berberich, Shady Elbassuoni, and Gerhard Weikum. 2013. Robust question answering over the web of linked data. In *CIKM*.