

BREAK It Down: A Question Understanding Benchmark

Tomer Wolfson^{1,3}, Mor Geva^{1,3}, Ankit Gupta¹, Matt Gardner³,
Yoav Goldberg^{2,3}, Daniel Deutch¹, Jonathan Berant^{1,3}

¹Tel Aviv University ²Bar-Ilan University ³Allen Institute for AI
{tomerwol,morgeva}@mail.tau.ac.il, ankitgupta.iitkanpur@gmail.com,
mattg@allenai.org, yoav.goldberg@gmail.com,
danielde@post.tau.ac.il, joberant@cs.tau.ac.il

Abstract

Understanding natural language questions entails the ability to break down a question into the requisite steps for computing its answer. In this work, we introduce a Question Decomposition Meaning Representation (QDMR) for questions. QDMR constitutes the ordered list of steps, expressed through natural language, that are necessary for answering a question. We develop a crowdsourcing pipeline, showing that quality QDMRs can be annotated at scale, and release the **BREAK** dataset, containing over 83K pairs of questions and their QDMRs. We demonstrate the utility of QDMR by showing that (a) it can be used to improve open-domain question answering on the **HOTPOTQA** dataset, (b) it can be deterministically converted to a pseudo-SQL formal language, which can alleviate annotation in semantic parsing applications. Last, we use **BREAK** to train a sequence-to-sequence model with copying that parses questions into QDMR structures, and show that it substantially outperforms several natural baselines.

1 Introduction

Recently, increasing work has been devoted to models that can reason and integrate information from multiple parts of an input. This includes reasoning over images (Antol et al., 2015; Johnson et al., 2017; Suhr et al., 2019; Hudson and Manning, 2019), paragraphs (Dua et al., 2019), documents (Welbl et al., 2018; Talmor and Berant, 2018; Yang et al., 2018), tables (Pasupat and Liang, 2015), and more. *Question answering* (QA) is commonly

used to test the ability to reason, where a complex natural language question is posed, and is to be answered given a particular context (text, image, etc.). Although questions often share structure across tasks and modalities, understanding the language of complex questions has thus far been addressed within each task in isolation. Consider the questions in Figure 1, all of which express operations such as fact chaining and counting. Additionally, humans can take a complex question and break it down into a sequence of simpler questions even when they are unaware of what or where the answer is. This ability, to compose and decompose questions, lies at the heart of human language (Pelletier, 1994) and allows us to tackle previously unseen problems. Thus, better question understanding models should improve performance and generalization in tasks that require multi-step reasoning or that do not have access to substantial amounts of data.

In this work we propose *question understanding* as a standalone language understanding task. We introduce a formalism for representing the meaning of questions that relies on *question decomposition*, and is agnostic to the information source. Our formalism, Question Decomposition Meaning Representation (QDMR), is inspired by database query languages (SQL; SPARQL), and by semantic parsing (Zelle and Mooney, 1996; Zettlemoyer and Collins, 2005; Clarke et al., 2010), in which questions are given full meaning representations.

We express complex questions via simple (“atomic”) questions that can be executed in sequence to answer the original question. Each atomic question can be mapped into a small set of *formal operations*, where each operation either selects a set of entities, retrieves information about

their attributes, or aggregates information over entities. While this has been formalized in knowledge-base (KB) query languages (Chamberlin and Boyce, 1974), the same intuition can be applied to other modalities, such as images and text. QDMR abstracts away the context needed to answer the question, allowing in principle to query multiple sources for the same question.

In contrast to semantic parsing, QDMR operations are expressed through *natural language*, facilitating annotation at scale by non-experts. Figure 1 presents examples of complex questions on three different modalities. The middle box lists the natural language decompositions provided for each question, and the bottom box displays their corresponding formal queries.

QDMR serves as the formalism for creating BREAK, a question decomposition dataset of 83,978 questions over ten datasets and three modalities. BREAK is collected via crowdsourcing, with a user interface that allows us to train crowd-workers to produce quality decompositions (§3). Validating the quality of annotated structures reveals 97.4% to be correct (§4).

We demonstrate the utility of QDMR in two setups. First, we regard the task of open-domain QA over multi-hop questions from the HOTPOTQA dataset. Combining QDMR structures in BREAK with a reading comprehension (RC) model (Min et al., 2019b) improves F_1 from 43.3 to 52.4 (§5). Second, we show that decompositions in BREAK possess high annotation consistency, which indicates that annotators produce high-quality QDMRs (§4.3). In §6 we discuss how these QDMRs can be used as a strong proxy for full logical forms in semantic parsing.

We use BREAK to train a neural QDMR parser that maps questions into QDMR representations, based on a sequence-to-sequence model with copying (Gu et al., 2016). Manual analysis of generated structures reveals an accuracy of 54%, showing that automatic QDMR parsing is possible, though still far from human performance (§7).

To conclude, our contributions are:

- Proposing the task of question understanding and introducing the QDMR formalism for representing the meaning of questions (§2)
- The BREAK dataset, which consists of 83,978 examples sampled from 10 datasets over three distinct information sources (§3)

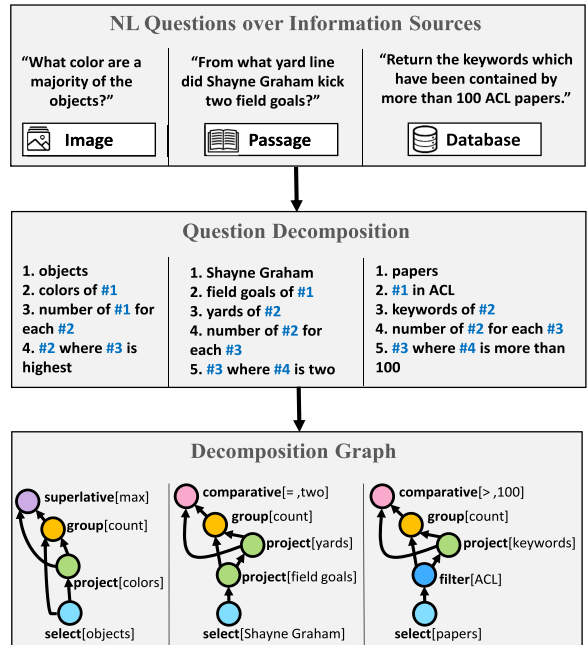


Figure 1: Questions over different sources share a similar compositional structure. Natural language questions from multiple sources (top) are annotated with the QDMR formalism (middle) and deterministically mapped into a pseudo-formal language (bottom).

- Showing how QDMR can be used to improve open-domain question answering (§5), as well as alleviate the burden of annotating logical forms in semantic parsing (§6)
- A QDMR parser based on a sequence-to-sequence model with copying mechanism (§7)

The BREAK dataset, models, and entire codebase are publicly available at: <https://github.com/tomerwolgithub/Break>.

2 Question Decomposition Formalism

In this section we define the QDMR formalism for domain agnostic question decomposition.

QDMR is primarily inspired by SQL (Codd, 1970; Chamberlin and Boyce, 1974). However, while SQL was designed for relational databases, QDMR also aims to capture the meaning of questions over unstructured sources such as text and images. Thus, our formalism abstracts away from SQL by assuming an underlying “idealized” KB, which contains all entities and relations expressed in the question. This abstraction enables QDMR to be unrestricted to a particular modality, with its operators to be executed also against text

Operator	Template / Signature	Question	Decomposition
Select	Return [entities] $w \rightarrow S_e$	How many touchdowns were scored overall?	1. Return touchdowns 2. Return the number of #1
Filter	Return [ref] [condition] $S_o, w \rightarrow S_o$	I would like a flight from Toronto to San Diego please.	1. Return flights 2. Return #1 from Toronto 3. Return #2 to San Diego
Project	Return [relation] of [ref] $w, S_e \rightarrow S_o$	Who is the head coach of the Los Angeles Lakers?	1. Return the Los Angeles Lakers 2. Return the head coach of #1
Aggregate	Return [aggregate] of [ref] $w_{agg}, S_o \rightarrow n$	How many states border Colorado?	1. Return Colorado 2. Return border states of #1 3. Return the number of #2
Group	Return [aggregate] [ref1] for each [ref2] $w_{agg}, S_o, S_e \rightarrow S_n$	How many female students are there in each club?	1. Return clubs 2. Return female students of #1 3. Return the number of #2 for each #1
Superlative	Return [ref1] where [ref2] is [highest / lowest] $S_e, S_n, w_{sup} \rightarrow S_e$	What is the keyword, which has been contained by the most number of papers?	1. Return papers 2. Return keywords of #1 3. Return the number of #1 for each #2 4. Return #2 where #3 is highest
Comparative	Return [ref1] where [ref2] [comparison] [number] $S_e, S_n, w_{com}, n \rightarrow S_e$	Who are the authors who have more than 500 papers?	1. Return authors 2. Return papers of #1 3. Return the number of #2 for each of #1 4. Return #1 where #3 is more than 500
Union	Return [ref1] , [ref2] $S_o, S_o \rightarrow S_o$	Tell me who the president and vice-president are?	1. Return the president 2. Return the vice-president 3. Return #1, #2
Intersection	Return [relation] in both [ref1] and [ref2] $w, S_e, S_e \rightarrow S_o$	Show the parties that have representatives in both New York state and representatives in Pennsylvania state.	1. Return representatives 2. Return #1 in New York state 3. Return #1 in Pennsylvania state 4. Return parties in both #2 and #3
Discard	Return [ref1] besides [ref2] $S_o, S_o \rightarrow S_o$	Find the professors who are not playing Canoeing.	1. Return professors 2. Return #1 playing Canoeing 3. Return #1 besides #2
Sort	Return [ref1] sorted by [ref2] $S_e, S_n \rightarrow \langle e_1 \dots e_k \rangle$	Find all information about student addresses, and sort by monthly rental.	1. Return students 2. Return addresses of #1 3. Return monthly rental of #2 4. Return #2 sorted by #3
Boolean	Return [if / is] [ref1] [condition] [ref2] $S_o, w, S_o \rightarrow b$	Were Scott Derrickson and Ed Wood of the same nationality?	... 3. Return the nationality of #1 4. Return the nationality of #2 5. Return if #3 is the same as #4
Arithmetic	Return the [arithmetic] of [ref1] and [ref2] $w_{ari}, n, n \rightarrow n$	How many more red objects are there than blue objects?	... 3. Return the number of #1 4. Return the number of #2 5. Return the difference of #3 and #4

Table 1: The 13 operator types of QDMR steps. Listed are, the natural language template used to express the operator, the operator signature, and an example question that uses the query operator in its decomposition.

and images, while allowing in principle to query multiple modalities for the same question.¹

QDMR Definition Given a question x , its QDMR is a sequence of n steps, $\mathbf{s} = \langle s^1, \dots, s^n \rangle$, where each step s^i corresponds to a single query

¹A system could potentially answer “Name the political parties of the most densely populated country”, by retrieving “the most densely populated country” using a database query, and “the political parties of #1” via an RC model.

operator f^i (see Table 1). A step, s^i is a sequence of tokens, $s^i = (s_1^i, \dots, s_{m_i}^i)$, where a token s_k^i is either a word from a predefined lexicon L_x (details in §3) or a *reference token*, referring to the result of a previous step s^j , where $j < i$. The last step, s^n returns the answer to x .

Decomposition Graph QDMR structures can be represented as a directed acyclic graph (DAG), used for evaluating QDMR parsing models (§7.1).

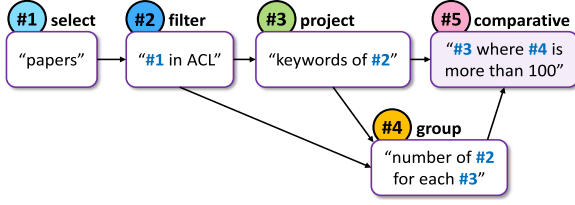


Figure 2: QDMR of the question ‘‘Return the keywords which have been contained by more than 100 ACL papers.’’, represented as a decomposition graph.

Given QDMR, $\mathbf{s} = \langle s^1, \dots, s^n \rangle$, each step s^i is a node in the graph, labeled by its sequence of tokens and index i . Edges in the graph are induced by *reference tokens* to previous steps. Node s^i is connected by an incoming edge (s^j, s^i) , if $ref[s^j] \in (s^i, \dots, s^i_{m_i})$. That is, if one of the tokens in s^i is a reference to s_j . Figure 2 displays a sequence of QDMR steps, represented as a DAG.

QDMR Operators A QDMR step corresponds to one of 13 query operators. We designed the operators to be expressive enough to represent the meaning of questions from a diverse set of datasets (§3). QDMR assumes an underlying KB, \mathcal{K} , which contains all of the entities and relations expressed in its steps. A relation, r , is a function mapping two arguments to whether r holds in \mathcal{K} : $\llbracket r(x, y) \rrbracket_{\mathcal{K}} \in \{\text{true}, \text{false}\}$. The operators operate over: (i) sets of objects S_o , where *objects* o , are either *numbers* n , *boolean* values b , or *entities* e in \mathcal{K} ; (ii) a closed set of phrases w_{op} , describing logical operations; and (iii) natural language phrases w , representing entities and relations in \mathcal{K} . We assume the existence of grounding functions that map a phrase w to concrete constants in \mathcal{K} . Table 2 describes the aforementioned constructs. In addition, we define the function $\text{map}_{\mathcal{K}}(S_e, S_o)$ which maps entity $e \in S_e$ to the set of corresponding objects from S_o . Each $o \in S_o$ corresponds to an $e \in S_e$ by being contained in the result of a sequence of PROJECT and GROUP operations applied to e :²

$$\text{map}_{\mathcal{K}}(S_e, S_o) = \{(e, o) \mid e \in S_e, o \in S_o, \\ o \in \text{op}_k \circ \dots \circ \text{op}_1(e)\}.$$

We now formally define each QDMR operator and provide concrete examples in Table 1.

- **SELECT**: Computes the set of entities in \mathcal{K} corresponding to w : $\text{select}(w) = \text{ground}_{\mathcal{K}}^e(w)$.

²The sequence of operations $\text{op}_1, \dots, \text{op}_k$ is traced using the references to previous steps in the QDMR structure.

Function	Description
agg	Given a phrase w_{agg} which describes an aggregate operation, agg denotes the corresponding operation. Either max , min , count , sum or avg .
sup	Given w_{sup} describing a superlative, it denotes the corresponding function. Either arg max or arg min .
com	Given w_{com} describing a comparison, it denotes the corresponding relation out of: $<$, \leq , $>$, \geq , $=$, \neq .
ari	Given w_{ari} describing an arithmetic operation, it denotes the corresponding operation out of: $+$, $-$, $*$, $/$.
$\text{ground}_{\mathcal{K}}^e(w)$	Given a natural language phrase w , it returns the set of corresponding KB entities, S_e .
$\text{ground}_{\mathcal{K}}^r(w)$	Given a natural language phrase w , it returns the corresponding KB relation, r .

Table 2: Functions used for grounding natural language phrases in numerical operators or KB entities.

- **FILTER**: Filters a set of objects so that it follows the condition expressed by w :

$$\text{filter}(S_o, w) = S_o \cap \{o \mid \llbracket r(e, o) \rrbracket_{\mathcal{K}} \equiv \text{true}\},$$

$$\text{where } r = \text{ground}_{\mathcal{K}}^r(w), e = \text{ground}_{\mathcal{K}}^e(w).$$

- **PROJECT**: Computes the objects that relate to input entities S_e with the relation expressed by w ,

$$\text{proj}(w, S_e) = \{o \mid \llbracket r(e, o) \rrbracket_{\mathcal{K}} \\ \equiv \text{true}, e \in S_e\},$$

$$\text{where } r = \text{ground}_{\mathcal{K}}^r(w).$$

- **AGGREGATE**: The result of applying an aggregate operation: $\text{aggregate}(w_{agg}, S_o) = \{\text{agg}(S_o)\}$.
- **GROUP**: Receives a set of ‘‘keys’’, S_e , and a set of corresponding ‘‘values’’, S_o . It outputs a set of numbers, each corresponding to a key $e \in S_e$. Each number results from applying aggregate , w_{agg} to the subset of values corresponding to e .

$$\text{group}(w_{agg}, S_o, S_e) = \\ \{\text{agg}(V_o(e)) \mid e \in S_e\},$$

$$\text{where } V_o(e) = \{o \mid \langle e, o \rangle \in \text{map}_{\mathcal{K}}(S_e, S_o)\}.$$

- **SUPERLATIVE**: Receives entity set S_e and number set S_n . Each number $n \in S_n$ is the result of a mapping from an entity $e \in S_e$. It returns a subset of S_e for which the corresponding number is either highest/lowest as indicated by w_{sup} .

$$\text{super}(S_e, S_n, w_{sup}) = \\ \{\text{sup}(\text{map}_{\mathcal{K}}(S_e, S_n))\}.$$

- **COMPARATIVE:** Receives entity set S_e and number set S_n . Each $n \in S_n$ is the result of a mapping from an $e \in S_e$. It returns a subset of S_e for which the comparison with n' , represented by w_{com} , holds.

$$\text{comparative}(S_e, S_n, w_{com}, n') = \{e \mid \langle e, n \rangle \in \text{map}_{\mathcal{K}}(S_e, S_n), \text{com}(n, n') \equiv \text{true}\}.$$

- **UNION:** Denotes the union of object sets: $\text{union}(S_o^1, S_o^2) = S_o^1 \cup S_o^2$.
- **DISCARD:** Denotes the set difference of two object sets: $\text{discard}(S_o^1, S_o^2) = S_o^1 \setminus S_o^2$.
- **INTERSECTION:** Computes the intersection of its entity sets and returns all objects which relate to the entities with the relation expressed by w .

$$\text{intersect}(w, S_e^1, S_e^2) = \{o \mid e \in S_e^1 \cap S_e^2, \llbracket r(e, o) \rrbracket_{\mathcal{K}} \equiv \text{true}, r = \text{ground}_{\mathcal{K}}^r(w)\}.$$

- **SORT:** Orders a set of entities according to a corresponding set of numbers. Each number n_i is the result of a mapping from entity e_i .

$$\text{sort}(S_e, S_n) = \{\langle e_{i_1} \dots e_{i_m} \rangle \mid \langle e_{i_j}, n_{i_j} \rangle \in \text{map}_{\mathcal{K}}(S_e, S_n), n_{i_1} \leq \dots \leq n_{i_m}\}.$$

- **BOOLEAN:** Returns whether the relation expressed by w holds between the input objects: $\text{boolean}(S_o^1, w, S_o^2) = \{\llbracket r(o_1, o_2) \rrbracket_{\mathcal{K}}\}$, where $r = \text{ground}_{\mathcal{K}}^r(w)$ and S_o^1, S_o^2 are singleton sets containing o_1, o_2 respectively.
- **ARITHMETIC:** Computes the application of an arithmetic operation: $\text{arith}(w_{ari}, S_n^1, S_n^2) = \{\text{ari}(n_1, n_2)\}$, where S_n^1, S_n^2 are singleton sets containing n_1, n_2 respectively.

High-level Decompositions In QDMR, each step corresponds to a single logical operator. In certain contexts, a less granular decomposition might be desirable, where sub-structures containing multiple operators could be collapsed to a single node. This can be easily achieved in QDMR by merging certain adjacent nodes in its DAG structure. When examining existing RC datasets (Yang et al., 2018; Dua et al., 2019), we observed that long spans in the question often match long spans in the text, due to existing practices of generating questions

Question: “The actress that played Pearl Gallagher on the TV sitcom *Different Strokes* is also the voice of an animated character that debuted in what year?”

High-level QDMR:

1. Return the actress that played Pearl Gallagher on the TV sitcom *Different Strokes*
 2. Return the animated character that #1 is the voice of
 3. Return the year that #2 debuted in
-

Figure 3: Example of a *high-level* QDMR. Step #1 merges together SELECT and multiple FILTER steps.

via crowdsourcing. In such cases, decomposing the long spans into multiple steps and having an RC model process each step independently, increases the probability of error. Thus, to promote the usefulness of QDMR for current RC datasets and models, we introduce *high-level* QDMR, by merging the following operators:

- **SELECT + PROJECT on named entities:** For the question, “What is the birthdate of Jane?” its *high-level* QDMR would be “return the birthdate of Jane” as opposed to the more granular, “return Jane; return birthdate of #1”.
- **SELECT + FILTER:** Consider the first step of the example in Figure 3. It contains both a SELECT operator (“return actress”) as well as two FILTER conditions (“that played...”, “on the TV sitcom...”).
- **FILTER + GROUP + COMPARATIVE:** Certain *high-level* FILTER steps contain implicit grouping and comparison operations. E.g., “return yard line scores in the fourth quarter; return #1 that both teams scored from”. Step #2 contains an implicit GROUP of team per yard line and a COMPARATIVE returning the lines where exactly two teams scored.

We provide both granular and high-level QDMRs for a random subset of RC questions (see Table 3). The concrete utility of *high-level* QDMR to open-domain QA is presented in §5.

3 Data Collection

Our annotation pipeline for generating BREAK consisted of three phases. First, we collected complex questions from existing QA benchmarks. Second, we crowdsourced the QDMR annotation of these questions. Finally, we validated worker annotations in order to maintain their quality.

Dataset	Example	Original	BREAK
ACADEMIC (DB)	Return me the total citations of all the papers in the VLDB conference.	195	195
ATIS (DB)	What is the first flight from Atlanta to Baltimore that serves lunch?	5,283	4,906
GEOQUERY (DB)	How high is the highest point in the largest state?	880	877
SPIDER (DB)	How many transactions correspond to each invoice number?	10,181	7,982
CLEVR-HUMANS (IMAGES)	What is the number of cylinders divided by the number of cubes?	32,164	13,935
NLVR2 (IMAGES)	If there are only two dogs pulling one of the sleds?	29,680	13,517
COMQA (TEXT)	What was Gandhi's occupation before becoming a freedom fighter?	11,214	5,520
CWQ (TEXT)	Robert E Jordan is part of the organization started by whom?	34,689	2,988, 2,991 ^{high}
DROP (TEXT)	Approximately how many years did the churches built in 1909 survive?	96,567	10,230, 10,262 ^{high}
HOTPOTQA-HARD (TEXT)	Benjamin Halfpenny was a footballer for a club that plays its home matches where?	23,066	10,575 ^{high}
Total:			83,978

Table 3: The QA datasets in BREAK. Lists the number of examples in the original dataset and in BREAK. Numbers of *high-level* QDMRs are denoted by *high*.

Question Collection Questions in BREAK were randomly sampled from ten QA datasets over the following tasks (Table 3):

- **Semantic Parsing:** Mapping natural language utterances into formal queries, to be executed on a target KB (Price, 1990; Zelle and Mooney, 1996; Li and Jagadish, 2014; Yu et al., 2018).
- **Reading Comprehension (RC):** Questions that require understanding of a text passage by reasoning over multiple sentences (Talmor and Berant, 2018; Yang et al., 2018; Dua et al., 2019; Abujabal et al., 2019).
- **Visual Question Answering (VQA):** Questions over images that require both visual and numerical reasoning skills (Johnson et al., 2017; Suhr et al., 2019).

All questions collected were composed by human annotators.³ HOTPOTQA questions were all sampled from the *hard* split of the dataset.

QDMR Annotation A key question is whether it is possible to train non-expert annotators to produce high-quality QDMRs. We designed an annotation interface (Figure 4), where workers are first given explanations and examples on how to identify and phrase each of the operators in Table 1. Then, workers decompose questions into a list of

³Except for COMPLEXWEBQUESTIONS (CWQ), where annotators paraphrased automatically generated questions.

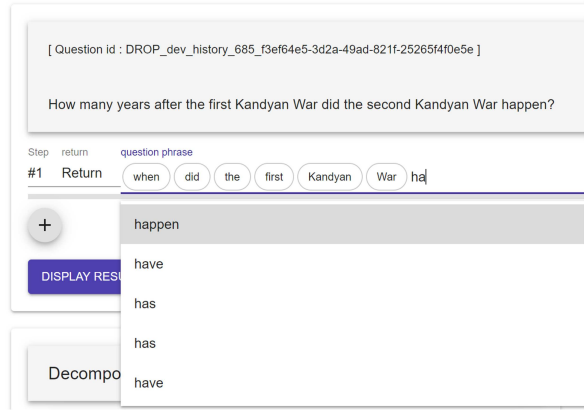


Figure 4: User interface for decomposing a complex question that uses a closed lexicon of tokens.

steps, where they are only allowed to use words from a lexicon L_x , which contains: (a) words appearing in the question (or their automatically computed inflections), (b) words from a small pre-defined list of 66 function word such as, ‘*if*’, ‘*on*’, ‘*for each*’, or (c) *reference tokens* that refer to the results of a previous step. This ensures that the language used by workers is consistent across examples, while being expressive enough for the decomposition. Our annotation interface presents workers with the question only, so they are agnostic to the original modality of the question. The efficacy of this process is explored in §4.2.

We used Amazon Mechanical Turk to crowd-source QDMR annotation. In each task, workers decomposed questions, paying them \$0.40 per question, which amounts to an average pay of \$12 per hour. Overall, we collected 83,978 examples using 64 distinct workers. The dataset was partitioned into train/development/test sets following the partitions in the original datasets. During partition, we made sure that development and test samples do not share the same context.

Worker Validation To ensure worker quality, we initially published qualification tasks, open to all workers in the United States. The task required workers to carefully review the annotation instructions and decompose 10 example questions. The examples were selected so that each QDMR operation should appear in at least one of their decompositions (Table 1). In total, 64 workers were able to correctly decompose at least 8 examples and were qualified as annotators. To validate worker performance over time, we conducted random validations of annotations. Over 9K annotations were

Operator	QDMR	QDMR _{high}
SELECT	100%	100%
PROJECT	69.0%	35.6%
FILTER	53.2%	15.3%
AGGREGATE	38.1%	22.3%
BOOLEAN	30.0%	4.6%
COMPARATIVE	17.0%	1.0%
GROUP	9.7%	0.7%
SUPERLATIVE	6.3%	13.0%
UNION	5.5%	0.5%
ARITHMETIC	5.4%	11.2%
DISCARD	3.2%	1.2%
INTERSECTION	2.7%	2.8%
SORT	0.9%	0.0%
Total	60,150	23,828

Table 4: Operator prevalence in BREAK. Lists the percentage of QDMRs where the operator appears.

reviewed by experts throughout the annotation process. Only workers who consistently produced correct QDMRs for at least 90% of their tasks were allowed to continue as annotators.

4 Dataset Analysis

This section examines the properties of collected QDMRs in BREAK and analyzes their quality.

4.1 Quantitative Analysis

Overall, BREAK contains 83,978 decompositions, including 60,150 QDMRs and 23,828 examples with *high-level* QDMRs, which are exclusive to text modalities. Table 3 shows that data is proportionately distributed between questions over structured (DB) and unstructured modalities (text, images).

The distribution of QDMR operators is presented in Table 4, detailing the prevalence of each query operator⁴ (we automatically compute this distribution, as explained in §4.3). SELECT and PROJECT are the most common operators. Additionally, at least 10% of QDMRs contain operators such as GROUP and COMPARATIVE, which entail complex reasoning, in contrast to *high-level* QDMRs, where such operations are rare. This distinction sheds light on the reasoning types required for answering RC datasets (*high-level* QDMR) compared with more structured tasks (QDMR).

Table 5 details the distribution of QDMR sequence length. Most decompositions in QDMR

⁴Regarding the three merged operators of *high-level* QDMRs (§2), the first two operators are treated as SELECT, while the third is considered a FILTER.

Steps	QDMR	QDMR _{high}
1–2	10.7%	59.8%
3–4	44.9%	31.6%
5–6	27.0%	7.9%
7–8	10.1%	0.6%
9+	7.4%	0.2%

Table 5: The distribution over QDMR sequence length.

include 3–6 steps, whereas *high-level* QDMRs are much shorter, as a single SELECT often finds an entity described by a long noun phrase (see §2).

4.2 Quality Analysis

We describe the process of estimating the *correctness* of collected QDMR annotations. Similar to previous works (Yu et al., 2018; Kwiatkowski et al., 2019) we use expert judgments, where the experts had prepared the guidelines for the annotation task. Given a question and its annotated QDMR, (q, s) the expert determines the correctness of s using one of the following categories:

- Correct (\mathcal{C}): If s constitutes a list of QDMR operations that lead to correctly answering q .
- Granular (\mathcal{C}_G): If s is correct and none of its operators can be further decomposed.⁵
- Incorrect (\mathcal{I}): If s is in neither \mathcal{C} nor \mathcal{C}_G .

Examples of these expert judgments are shown in Figure 5. To estimate expert judgment of correctness, we manually reviewed a random sample of 500 QDMRs from BREAK. We classified 93.8% of the samples in \mathcal{C}_G and another 3.6% in \mathcal{C} . Thus, 97.4% of the samples constitute a correct decomposition of the original question. Workers have somewhat struggled with decomposing superlatives (e.g., “*biggest sphere*”), as evident from the first question in Figure 5. Collected QDMRs displayed similar estimates of \mathcal{C} , \mathcal{C}_G , and \mathcal{I} , regardless of their modality (DB, text, or image).

4.3 Annotation Consistency

As QDMR is expressed using natural language, it introduces variability into its annotations. We wish to validate the *consistency* of collected QDMRs, that is, whether we can correctly infer the formal

⁵For *high-level* QDMRs, the merged operators (§2) are considered to be fully decomposed.

Question 1: “What color is the biggest sphere in the picture?”

QDMR: (1) Return spheres; (2) Return #1 that is the biggest; (3) Return color of #2.

Expert Judgement: \mathcal{C} . Correct, but not fully decomposed. Step #2 should be broken down to: (2) Return size of #1; (3) Return #1 where #3 is highest.

Question 2: “What is the full name of each car maker, along with its id and how many models it produces?”

QDMR: (1) Return car makers; (2) Return models of #1; (3) Return number of #2 for each #1; (4) Return full names of #1; (5) Return ids of #1; (6) Return #4, #5, #3.

Expert Judgement: \mathcal{C}_G . Correct and fully decomposed.

Question 3: “Show the names and locations of institutions that are founded after 1990 and have the type Private.”

QDMR: (1) Return institutions; (2) Return #1 founded after 1990; (3) Return types of #1; (4) Return #1 where #3 is Private; (5) Return #2, #4; (6) Return names of #5; (7) Return locations of #5; (8) return #6, #7.

Expert Judgement: \mathcal{I} . Incorrect, as step #5 returns institutions that were either founded after 1990, or are Private.

Figure 5: Examples and justifications of expert judgment on collected QDMRs in BREAK.

QDMR operator (f^i) and its arguments from each step (s^i). To infer these formal representations, we developed an algorithm that goes over the QDMR structure step-by-step, and for each step s^i , uses a set of predefined templates to identify f^i and its arguments, expressed in s^i . This results in an execution graph (Figure 2), where the execution result of a parent node serves as input to its child. Figure 1 presents three QDMR decompositions along with the formal graphs output by our algorithm (lower box). Each node lists its operator (e.g., GROUP), its *constant input* listed in brackets (e.g., count) and its *dynamic input*, which are the execution results of its parent nodes.

Overall, 99.5% of QDMRs had all their steps mapped into pseudo-logical forms by our algorithm. To evaluate the correctness of the mapping algorithm, we randomly sampled 350 logical forms, and examined the structure of the formulas, assuming that words copied from the question correspond to entities and relations in an idealized KB (see §2). Of this sample, 99.4% of its examples had all of their steps, s^i , correctly mapped to the corresponding f^i . Overall, 93.1% of the examples were of fully accurate logical forms, with errors being due to QDMRs that were either incorrect or not fully decomposed (\mathcal{I} , \mathcal{C} in §4.2). Thus, a rule-based algorithm can map more than 93% of the annotations into a correct formal representation. This shows that our annotators produced consistent and high-quality QDMRs. Moreover, it suggests that non-experts can annotate questions with pseudo-logical forms, which can be used as

Algorithm 1 BREAKRC

```
1: procedure BREAKRC( $s$ : QDMR)
2:    $ansrs \leftarrow []$ 
3:   for  $s^i$  in  $s = \langle s^1, \dots, s^n \rangle$  do
4:      $op \leftarrow \text{OPTYPE}(s^i)$ 
5:      $refs \leftarrow \text{REFERENCEDSTEPS}(s^i)$ 
6:     if  $op$  is SELECT then
7:        $ans \leftarrow \text{ANSWER}(s^i)$ 
8:     else if  $op$  is FILTER then
9:        $\hat{s}^i \leftarrow \text{EXTRACTQUESTION}(s^i)$ 
10:       $ans_{tmp} \leftarrow \text{ANSWER}(\hat{s}^i)$ 
11:       $ans \leftarrow \text{INTERSECT}(ans_{tmp}, ansrs[refs[0]])$ 
12:     else if  $op$  is COMPARISON then
13:        $ans \leftarrow \text{COMPARESTEPS}(refs, s)$ 
14:     else  $\triangleright op$  is PROJECT
15:        $\hat{s}^i \leftarrow \text{SUBSTITUTEREF}(s^i, ansrs[refs[0]])$ 
16:        $ans \leftarrow \text{ANSWER}(\hat{s}^i)$ 
17:        $ansrs[i] \leftarrow ans$ 
18:   return  $ansrs[n]$ 
```

a cheap intermediate representation for semantic parsers (Yih et al., 2016), further discussed in §6.

5 QDMR for Open-domain QA

A natural setup for QDMR is in answering complex questions that require multiple reasoning steps. We compare models that exploit question decompositions to baselines that do not. We use the open-domain QA (“full-wiki”) setting of the HOTPOTQA dataset (Yang et al., 2018): Given a question, the QA model retrieves the relevant Wikipedia paragraphs and answers the question using these paragraphs.

5.1 Experimental Setup

We compare BREAKRC, a model that utilizes question decomposition to BERTQA, a standard QA model, based on BERT (Devlin et al., 2019), and present COMBINED, an approach that enjoys the benefits of both models.

BREAKRC Algorithm 1 describes the BREAKRC model, which uses high-level QDMR structures for answering open-domain multi-hop questions. We assume access to an Information Retrieval (IR) model and an RC model, and denote by ANSWER(\cdot) a function that takes a question as input, runs the IR model to obtain paragraphs, and then feeds those paragraphs as context for an RC model that returns a distribution over answers.

Given an input QDMR, $s = \langle s^1, \dots, s^n \rangle$, iterate over s step-by-step and perform the following. First, we extract the operation (line 4) and the previous steps referenced by s^i (line 5). Then, we compute the answer to s^i conditioned on the

extracted operator. For SELECT steps, we simply run the ANSWER(\cdot) function. For PROJECT steps, we substitute the reference to the previous step in s^i with its already computed answer, and then run ANSWER(\cdot). For FILTER steps,⁶ we use a simple rule to extract a ‘‘normalized question’’, \hat{s}^i from s^i and get an intermediate answer ans_{tmp} with ANSWER(\hat{s}^i). We then ‘‘intersect’’ ans_{tmp} with the referenced answer by multiplying the probabilities provided by the RC model and normalizing. For COMPARISON steps, we compare, with a discrete operation, the numbers returned by the referenced steps. The final answer is the highest probability answer of step s^n .

As our IR model we use bigram TF-IDF, proposed by Chen et al. (2017). Because the RC model is run on single-hop questions, we use the BERT-based RC model from Min et al. (2019b), trained solely on SQuAD (Rajpurkar et al., 2016).

BERTQA Baseline AS BREAKRC exploits question decompositions, we compare it with a model that does not. BERTQA receives as input the original natural language question, x . It uses the same IR model as BREAKRC to retrieve paragraphs for x . For a fair comparison, we set its number of retrieved paragraphs such that it is identical to BREAKRC (namely, 10 paragraphs for each QDMR step that involves IR). Similar to BREAKRC, retrieved paragraphs are fed to a pretrained BERT-based RC model (Min et al., 2019b) to answer x . In contrast to BREAKRC, that is trained on SQuAD, BERTQA is trained on the target dataset (HOTPOTQA), giving it an advantage over BREAKRC.

A COMBINED Approach Last, we present an approach that combines the strengths of BREAKRC and BERTQA. In this approach, we use the QDMR decomposition to improve *retrieval* only. Given a question x and its QDMR s , we run BREAKRC on s , but in addition to storing *answers*, we also store all the *paragraphs* retrieved by the IR model. We then run BERTQA on the question x and the top-10 paragraphs retrieved by BREAKRC, sorted by their IR ranking. This approach resembles that of Qi et al. (2019).

The advantage of COMBINED is that we do not need to develop an answering procedure for each QDMR operator separately, which involves dif-

⁶INTERSECTION steps are handled in a manner similar to FILTER, but we omit the exact description for brevity.

Model	HOTPOTQA		
	EM	F ₁	IR
BERTQA	33.6	43.3	46.3
BREAKRC ^P	28.8	37.7	52.5
BREAKRC ^G	34.6	44.6	59.2
COMBINED ^P	38.3	49.3	52.5
COMBINED ^G	41.2	52.4	59.2
IR-NP	31.7	41.2	40.8
BREAKRC ^R	18.9	26.5	40.3
COMBINED ^R	32.7	42.6	40.3

Table 6: Open-domain QA results on HOTPOTQA.

ferent discrete operations such as comparison and intersection. Instead, we use BREAKRC to retrieve contexts, and an end-to-end approach to learn how to answer the question directly. This can often handle operators not implemented in BREAKRC, like BOOLEAN and UNION.

DATASET To evaluate our models, we use all 2,765 QDMR annotated examples of the HOTPOTQA development set found in BREAK. PROJECT and COMPARISON type questions account for 48% and 7% of examples respectively.

5.2 Results

Table 6 shows model performance on HOTPOTQA. We report EM and F₁ using the official HOTPOTQA evaluation script. IR measures the percentage of examples in which the IR model successfully retrieved *both* of the ‘‘gold paragraphs’’ necessary for answering the multi-hop question. To assess the potential utility of QDMR, we report results for BREAKRC^G, which uses *gold* QDMRs, and BREAKRC^P, which uses QDMRs predicted by a COPYNET parser (§7.2).

Retrieving paragraphs with decomposed questions substantially improves the IR metric from 46.3 to 59.2 (BREAKRC^G), or 52.5 (BREAKRC^P). This leads to substantial gains in EM and F₁ for COMBINED^G (43.3 to 52.4) and COMBINED^P (43.3 to 49.3). The EM and F₁ of BREAKRC^G are only slightly higher than BERTQA because BREAKRC does not handle certain operators, such as BOOLEAN steps (9.4% of the examples).

The majority of questions in HOTPOTQA combine SELECT operations with either PROJECT (also called ‘‘bridge’’ questions), COMPARISON, or FILTER. PROJECT and COMPARISON questions (Figure 6) were shown to be less susceptible

<p>Project Question: “The actor of Professor Sprout in Harry Potter acted as a mother in a sitcom. In that sitcom, what was the daughter’s name?”</p> <p>Gold Decomposition:</p> <ol style="list-style-type: none"> 1. Return actor of Professor Sprout In Harry Potter 2. Return sitcom that #1 acted as a mother in 3. Return daughter name in #2
<p>Comparison Question: “Which 1970’s film was released first, Charley and the Angel or The Boatniks?”</p> <p>Gold Decomposition:</p> <ol style="list-style-type: none"> 1. Return when Charley and the Angel was released 2. Return when The Boatniks was released 3. Return which is the lowest of #1 , #2

Figure 6: Examples of PROJECT and COMPARISON questions in HOTPOTQA (high-level QDMR).

Model	PROJECT			COMPARISON		
	EM	F ₁	IR	EM	F ₁	IR
BERTQA	22.8	31.0	31.6	42.9	51.7	75.8
BREAKRC ^P	25.4	33.7	52.9	34.7	50.4	68.9
BREAKRC ^G	32.2	41.9	59.8	44.5	57.6	78.0

Table 7: Results on PROJECT and COMPARISON questions from HOTPOTQA development set.

to reasoning shortcuts, i.e. they necessitate multi-step reasoning (Chen and Durrett, 2019; Jiang and Bansal, 2019; Min et al., 2019a). In Table 7 we report BREAKRC results on these question types, where it notably outperforms BERTQA.

Ablations In BREAKRC, multiple IR queries are issued, one at each step. To examine whether these multiple queries were the cause for performance gains, we built IR-NP, a model that issues multiple IR queries, one for each noun phrase in the question. Similar to COMBINED, the question and union of retrieved paragraphs are given as input to BERTQA. We observe that COMBINED substantially outperforms IR-NP, indicating that the structure of QDMR, rather than multiple IR queries, has led to improved performance.⁷

To test whether QDMR is better than a simple rule-based decomposition algorithm, we developed a model that decomposes a question by applying a set of predefined rules over the dependency tree of the question (full details in §7.2). COMBINED and BREAKRC were compared to COMBINED^R and BREAKRC^R, which use the rule-based decompositions. We observe that QDMR lead to substantially higher performance when compared to the rule-based decompositions.

⁷Issuing an IR query over each “content word” in the question, instead of each noun phrase, led to poor results.

6 QDMR for Semantic Parsing

As QDMR structures can be easily annotated at scale, a natural question is how far are they from fully executable queries (known to be expensive to annotate). As shown in §4.3, QDMRs can be mapped to *pseudo-logical forms* with high precision (93.1%) by extracting formal operators and arguments from their steps. The pseudo-logical form differs from an executable query in the lack of grounding of its arguments (entities and relations) in KB constants. This stems from the design of QDMR as a domain-agnostic meaning representation (§2). QDMR abstracts away from a concrete KB schema by assuming an underlying “idealized” KB, which contains all of its arguments.

Thus, QDMR can be viewed as an *intermediate representation* between a natural language question and an executable query. Such intermediate representations have already been discussed in prior work on semantic parsing. Kwiatkowski et al. (2013) and Choi et al. (2015) used *underspecified logical forms* as an intermediate representation. Guo et al. (2019) proposed a two-stage approach, separating between learning an intermediate text-to-SQL representation and the actual mapping to schema items. Works in the database community have particularly targeted the mapping of intermediate query representations into DB grounded queries, using schema mapping and join path inference (Androustopoulos et al., 1995; Li et al., 2014; Baik et al., 2019). We argue that QDMR can be used as an easy-to-annotate representation in such semantic parsers, bridging between natural language and full logical forms.

7 QDMR Parsing

We now present evaluation metrics and models for mapping questions into QDMR structures.

Task Definition Given a question x we wish to map it to its QDMR steps, $\mathbf{s} = \langle s^1, \dots, s^n \rangle$. One can frame this as a sequence-to-sequence problem where x is mapped to a string representing its decomposition. We add a special separating token $\langle \text{SEP} \rangle$, and define the target string to be $s_1^1, \dots, s_{m_1}^1, \langle \text{SEP} \rangle, s_1^2, \dots, s_{m_2}^2, \langle \text{SEP} \rangle, \dots, s_{m_n}^n$, where m_1, \dots, m_n are the number of tokens in each decomposition step.

Question: “Show me all the flights from Atlanta to Baltimore on any airline on Thursday”

Decomposition 1:

1. Return flights
2. Return #1 from Atlanta
3. Return #2 to Baltimore
4. Return #3 on Thursday
5. Return #4 from any airline

Decomposition 2:

1. Return flights from Atlanta to Baltimore
2. Return #1 on any airline
3. Return #2 on Thursday




Figure 8: Graph edit operations between the graphs of the two QDMRs in Figure 7.

Figure 7: Differences in granularity, step order, and wording between two decompositions.

7.1 Evaluation Metrics

We wish to assess the quality of a predicted QDMR, \hat{s} to a gold standard, s . Figure 7 lists various properties by which question decompositions may differ, such as *granularity* (e.g., steps 1–3 of decomposition 1 are merged into the first step of decomposition 2), *ordering* (e.g., the last two steps are swapped) and *wording* (e.g., using “from” instead of “on”). While such differences do not affect the overall semantics, the second decomposition can be further decomposed. To measure such variations, we introduce two types of evaluation metrics. *Sequence-based* metrics treat the decomposition as a sequence of tokens, applying standard text generation metrics. As such metrics ignore the QDMR graph structure, we also use *graph-based* metrics that compare the predicted graph $G_{\hat{s}}$ to the gold QDMR graph G_s (see §2).

Sequence-based scores, where higher values are better, are denoted by \uparrow . *Graph-based scores*, where lower values are better, are denoted by \downarrow .

- **Exact Match** \uparrow : Measures exact match between s and \hat{s} , either 0 or 1.
- **SARI** \uparrow (Xu et al., 2016): *SARI* is commonly used in tasks such as text simplification. Given s , we consider the sets of added, deleted, and kept n -grams when mapping the question x to s . We compute these three sets for both s and \hat{s} using the standard of up to 4-grams, then average (a) the F_1 for added n -grams between s and \hat{s} , (b) the F_1 for kept n -grams, and (c) the precision for the deleted n -grams.
- **Graph Edit Distance (GED)** \downarrow : A graph edit path is a sequence of node and edge edit operations (addition, deletion, and substitution), where each operation has a predefined

cost. *GED* computes the minimal-cost graph edit path required for transitioning from G_s to $G_{\hat{s}}$ (and vice versa), normalized by $\max(|G_s|, |G_{\hat{s}}|)$. Operation costs are 1 for insertion and deletion of nodes and edges. The substitution cost of two nodes u, v is set to be $1 - \text{Align}(u, v)$, where $\text{Align}(u, v)$ is the ratio of aligned tokens between these steps.

- **GED+** \downarrow : Comparing the QDMR graphs in Figure 8, we consider the splitting and merging of graph nodes. We implement *GED+*, a variant of *GED* with additional operations to merge (split) a set of nodes (node), based on the A* algorithm (Hart et al., 1968).⁸

7.2 QDMR Parsing Models

We present models for QDMR parsing, built over AllenNLP (Gardner et al., 2017).

- **COPY**: A model that copies the input question x , without introducing any modifications.
- **RULEBASED**: We defined 12 decomposition rules, to be applied over the dependency tree of the question, augmented with coreference relations. A rule is a regular expression over the question dependency tree, which invokes a decomposition operation when matched (Table 8). For example, the rule for relative clauses (`relcl`) breaks the question at the relative pronoun “that”, while adding a reference to the preceding part of the sentence. A full decomposition is obtained by recursively applying the rules until no rule is matched.
- **SEQ2SEQ**: A sequence-to-sequence neural model with a 5-layer LSTM encoder and attention at decoding time.

⁸Because of its exponential worst-case complexity, we compute *GED+* only for graphs with up to 5 nodes, covering 75.2% of the examples in the development set of BREAK.

Structure	Example
be-root	How many objects smaller than the matte object are silver [objects smaller than the matte object, How many #1 silver]
be-auxpass	Find the average rating star for each movie that are not reviewed by Brittany Harris. [Brittany Harris, the average rating star for each movie that not reviewed by #1]
do-subj	Year did the team with Baltimore Fight Song win the Superbowl? [team with Baltimore Fight Song, year did #1 win the Superbowl]
subj-do-have	Which team owned by Malcolm Glazer has Tim Howard playing? [team Tim Howard playing, #1 owned by Malcolm Glazer]
conjunction	Who trades with China and has a capital city called Khartoum? [Who has a capital city called Khartoum, #1 trades with China]
how-many	How many metallic objects appear in this image? [metallic objects appear in this image, the number of #1]
single-prep	Find the ids of the problems reported after 1978. [the problems reported after 1978, ids of #1]
multi-prep	what flights from Tacoma to Orlando on Saturday [flights, #1 from Tacoma, #2 to Orlando, #3 on Saturday]
relcl	Find all the songs that do not have a back vocal. [all the songs, #1 that do not have a back vocal]
superlative	What is the smallest state bordering ohio [state bordering ohio, the smallest #1]
acl-verb	Find the first names of students studying in 108. [students, #1 studying in 108, first names of #2]
sent-coref	Find the claim that has the largest total settlement amount. Return the effective date of the claim . [the claim that has the largest total settlement amount, the effective date of #1]

Table 8: The decomposition rules of RULEBASED. Rules are based on dependency labels, part-of-speech tags and coreference edges. Text fragments used for decomposition are in boldface.

- S2SDYNAMIC: SEQ2SEQ with a dynamic output vocabulary restricted to the closed set of tokens L_x available to crowd-workers (see §3).
- COPYNET: SEQ2SEQ with an added copy mechanism that allows copying tokens from the input sequence (Gu et al., 2016).

7.3 Results

Table 9 presents model performance on BREAK. Neural models outperform the RULEBASED baseline and perform reasonably well, with COPYNET obtaining the best scores across all metrics. This can be attributed to most of the tokens in a QDMR parse being copied from the original question.

Error Analysis To judge the quality of predicted QDMRs we sampled 100 predictions of COPYNET (Table 10) half of them being *high-level* QDMRs. For standard QDMR, 24% of the sampled predictions were an exact match, with an additional 30% being fully decomposed and semantically equivalent to the gold decompositions. For example, in the first row of Table 10, the gold decomposition first discards the number of cylinders, then counts the remaining objects. Instead, COPYNET opted to count both groups, then subtract the number of cylinders from the number of objects. This illustrates how different QDMRs may be equivalent.

For high-level examples (from RC datasets), as questions are often less structured, they require a deeper semantic understating from the decomposition model. Only 8% of the predictions were an exact match, with an additional 46% being semantically equivalent to the gold. The remaining 46% were of erroneous predictions (see Table 10).

8 Related Work

Question Decomposition Recent work on QA through question decomposition has focused mostly on single modalities (Gupta and Lewis, 2018; Guo et al., 2019; Min et al., 2019b). QA using neural modular networks has been suggested for both KBs and images by Andreas et al. (2016) and Hu et al. (2017). Question decomposition over text was proposed by Talmor and Berant (2018), however over a much more limited set of questions than in BREAK. Iyyer et al. (2017) have also decomposed questions to create a “sequential question answering” task. Their annotators viewed a web table and performed actions over it to retrieve the cells that constituted the answer. Conversely, we provided annotators only with the question, as QDMR is agnostic to the original context.

An opposite annotation cycle to ours was presented in Cheng et al. (2018). The authors generate sequences of simple questions which crowd-workers paraphrase into a compositional question. Questions in BREAK are composed by humans, and are then decomposed to QDMR.

Semantic Formalism Annotation Labeling corpora with a semantic formalism has often been reserved for expert annotators (Dahl et al., 1994; Zelle and Mooney, 1996; Abend and Rappoport, 2013; Yu et al., 2018). Recent work has focused on cheaply eliciting quality annotations from non-experts through crowdsourcing (He et al., 2016; Iyer et al., 2017; Michael et al., 2018). FitzGerald et al. (2018) facilitated non-expert annotation by introducing a formalism expressed in natural language for semantic-role-labeling. This mirrors QDMR, as both are expressed in natural language.

Relation to Other Formalisms QDMR is related to Dependency-based Compositional Semantics (Liang et al., 2013), as both focus on question representations. However, QDMR is designed

Data	Metric	COPY	RULEBASED	SEQ2SEQ	S2SDYNAMIC	COPYNET	COPYNET (test)
QDMR	<i>Exact Match</i> ↑	0.001	0.002	0.081	0.116	0.154	0.157
	<i>SARI</i> ↑	0.431	0.508	0.665	0.705	0.748	0.746
	<i>GED</i> ↓	0.937	0.799	0.398	0.363	0.318	0.322
	<i>GED+</i> ↓	1.813	1.722	1.424	1.137	0.941	0.984
QDMR _{high}	<i>Exact Match</i> ↑	0.001	0.010	0.001	0.015	0.081	0.083
	<i>SARI</i> ↑	0.501	0.554	0.379	0.504	0.722	0.722
	<i>GED</i> ↓	0.793	0.659	0.585	0.468	0.319	0.316
	<i>GED+</i> ↓	1.102	1.395	1.655	1.238	0.716	0.709

Table 9: Performance of QDMR parsing models on the development and test set. *GED+* is computed only for the subset of QDMR graphs with up to 5 nodes, covering 66.1% of QDMRs and 97.6% of high-level data.

Question	Gold	Prediction (COPYNET)	Analysis
“How many objects other than cylinders are there?”	(1) objects; (2) cylinders; (3) #1 besides #2; (4) number of #3.	(1) objects; (2) cylinders; (3) number of #1; (4) number of #2; (5) difference of #3 and #4.	sem. equiv. (30%)
“Where is the youngest teacher from?”	(1) teachers; (2) the youngest of #1; (3) where is #2 from.	(1) youngest teacher; (2) where is #1.	incorrect (46%)
“Kyle York is the Chief Strategy Officer of a company acquired by what corporation in 2016?”	(1) company that Kyle York is the Chief Strategy Officer of; (2) corporation that acquired #1 in 2016.	(1) company that Kyle York is the Chief Strategy Officer of; (2) corporation in 2016 that #1 was acquired by.	sem. equiv. (46%)
“Dayton’s Devils had a cameo from the ‘MASH’ star who played what role on the show?”	(1) MASH star that Dayton’s Devils had a cameo from; (2) role that #1 played on the show.	(1) the MASH that Dayton’s Devils had a cameo; (2) what role on the show star of #1 played.	incorrect (46%)

Table 10: Manual error analysis of the COPYNET model predictions. Lower examples are of *high-level* QDMRs.

to facilitate annotations, while Dependency-based Compositional Semantics is centered on paralleling syntax. Domain-independent intermediate representations for semantic parsers were proposed by Kwiatkowski et al. (2013) and Reddy et al. (2016). As there is no consensus on the ideal meaning representation for semantic parsing, representations are often chosen based on the particular execution setup: SQL is used for relational databases (Yu et al., 2018), SPARQL for graph KBs (Yih et al., 2016), while other ad-hoc languages are used based on the task at hand. We frame QDMR as an easy-to-annotate formalism that can be potentially converted to other representations, depending on the task. Last, AMR (Banarescu et al., 2013) is a meaning representation for sentences. Instead of representing general language, QDMR represents questions, which are important for QA systems, and for probing models for reasoning.

9 Conclusion

In this paper, we presented a formalism for question understanding. We have shown it is possible to train crowd-workers to produce such representations with high quality at scale, and created BREAK, a benchmark for question decomposition with over 83K decompositions of

questions from 10 datasets and 3 modalities (DB, images, text). We presented the utility of QDMR for both open-domain question answering and semantic parsing, and constructed a QDMR parser with reasonable performance. QDMR proposes a promising direction for modeling question understanding, which we believe will be useful for multiple tasks in which reasoning is probed through questions.

Acknowledgments

This work was completed in partial fulfillment for the PhD of Tomer Wolfson. This research was partially supported by The Israel Science Foundation (grants 942/16 and 978/17), and The Yandex Initiative for Machine Learning and the European Research Council (ERC) under the European Union Horizons 2020 research and innovation programme (grant ERC DELPHI 802800).

References

Omri Abend and Ari Rappoport. 2013. Universal conceptual cognitive annotation (UCCA). In *Association for Computational Linguistics (ACL)*.

- 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 *North American Association for Computational Linguistics (NAACL)*.
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016. Learning to compose neural networks for question answering. In *Human Language Technology and North American Association for Computational Linguistics (HLT/NAACL)*.
- Ion Androutsopoulos, Graeme D. Ritchie, and Peter Thanisch. 1995. Natural language interfaces to databases – an introduction. *Journal of Natural Language Engineering*, 1:29–81.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *International Conference on Computer Vision (ICCV)*, pages 2425–2433.
- Christopher Baik, Hosagrahar Visvesvaraya Jagadish, and Yunyao Li. 2019. Bridging the semantic gap with SQL query logs in natural language interfaces to databases. *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 374–385.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *7th Linguistic Annotation Workshop and Interoperability with Discourse*.
- Donald D. Chamberlin and Raymond F. Boyce. 1974. SEQUEL: A structured English query language. In *Proceedings of the 1974 ACM SIGFIDET (now SIGMOD) Workshop on Data Description, Access and Control*, pages 249–264. ACM.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Association for Computational Linguistics (ACL)*.
- Jifan Chen and Greg Durrett. 2019. Understanding dataset design choices for multi-hop reasoning. In *Association for Computational Linguistics (ACL)*.
- Jianpeng Cheng, Siva Reddy, and Mirella Lapata. 2018. Building a neural semantic parser from a domain ontology. *ArXiv*, abs/1812.10037. Version 1.
- Eunsol Choi, Tom Kwiatkowski, and Luke Zettlemoyer. 2015. Scalable semantic parsing with partial ontologies. In *Association for Computational Linguistics (ACL)*.
- James Clarke, Dan Goldwasser, Ming-Wei Chang, and Dan Roth. 2010. Driving semantic parsing from the world’s response. In *Computational Natural Language Learning (CoNLL)*, pages 18–27.
- Edgar F. Codd. 1970. A relational model of data for large shared data banks. *Communications of the ACM*, 13(6):377–387.
- Deborah A. Dahl, Madeleine Bates, Michael Brown, William M. Fisher, Kate Hunicke-Smith, David S. Pallett, Christine Pao, Alexander I. Rudnicky, and Elizabeth Shriberg. 1994. Expanding the scope of the ATIS task: The ATIS-3 corpus. In *Workshop on Human Language Technology*, pages 43–48.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Association for Computational Linguistics (NAACL)*.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *Human Language Technology and North American Association for Computational Linguistics (HLT/NAACL)*.
- Nicholas FitzGerald, Julian Michael, Luheng He, and Luke S. Zettlemoyer. 2018. Large-scale QA-SRL parsing. In *Association for Computational Linguistics (ACL)*.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Taffjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. 2017. In AllenNLP: A deep semantic natural language processing platform, arXiv, abs/1803.07640v2.

- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O. K. Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *Association for Computational Linguistics (ACL)*.
- Jiaqi Guo, Zecheng Zhan, Yan Gao, Yan Xiao, Jian-Guang Lou, Ting Liu, and Dongmei Zhang. 2019. Towards complex text-to-SQL in cross-domain database with intermediate representation. In *Association for Computational Linguistics (ACL)*.
- Nitish Gupta and Mike Lewis. 2018. Neural compositional denotational semantics for question answering. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107.
- Luheng He, Julian Michael, Mike Lewis, and Luke Zettlemoyer. 2016. Human-in-the-loop parsing. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. 2017. Learning to reason: End-to-end module networks for visual question answering. In *International Conference on Computer Vision (ICCV)*.
- Drew A. Hudson and Christopher D. Manning. 2019. GQA: A new dataset for real-world visual reasoning and compositional question answering. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Srini Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. 2017. Learning a neural semantic parser from user feedback. In *Association for Computational Linguistics (ACL)*.
- Mohit Iyyer, Wen-tau Yih, and Ming-Wei Chang. 2017. Search-based neural structured learning for sequential question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1821–1831.
- Yichen Jiang and Mohit Bansal. 2019. Avoiding reasoning shortcuts: Adversarial evaluation, training, and model development for multi-hop QA. In *Association for Computational Linguistics (ACL)*.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross B. Girshick. 2017. CleVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *Computer Vision and Pattern Recognition (CVPR)*.
- Tom Kwiatkowski, Eunsol Choi, Yoav Artzi, and Luke Zettlemoyer. 2013. Scaling semantic parsers with on-the-fly ontology matching. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Fei Li and Hosagrahar Visvesvaraya Jagadish. 2014. NaLIR: An interactive natural language interface for querying relational databases. In *International Conference on Management of Data, SIGMOD*.
- Fei Li, Tianyin Pan, and Hosagrahar Visvesvaraya Jagadish. 2014. Schema-free SQL. In *International Conference on Management of Data, SIGMOD*, pages 1051–1062.
- Percy Liang, Michael I. Jordan, and Dan Klein. 2013. Learning dependency-based compositional semantics. *Computational Linguistics*, 39:389–446.
- Julian Michael, Gabriel Stanovsky, Luheng He, Ido Dagan, and Luke Zettlemoyer. 2018. Crowdsourcing question–answer meaning representations. In *North American Association for Computational Linguistics (NAACL)*.
- Sewon Min, Eric Wallace, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019a. Compositional questions do not necessitate multi-hop reasoning. In *Association for Computational Linguistics (ACL)*.

- Sewon Min, Victor Zhong, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2019b. Multi-hop reading comprehension through question decomposition and rescoring. In *Association for Computational Linguistics (ACL)*.
- Panupong Pasupat and Percy Liang. 2015. Compositional semantic parsing on semi-structured tables. In *Association for Computational Linguistics (ACL)*.
- Francis Jeffrey Pelletier. 1994. The principle of semantic compositionality. *Topoi*, 13(1):11–24.
- P. J. Price. 1990. Evaluation of spoken language systems: The ATIS domain. In *Proceedings of the Third DARPA Speech and Natural Language Workshop*, pages 91–95.
- Peng Qi, Xiaowen Lin, Leo Mehr, Zijian Wang, and Christopher D. Manning. 2019. Answering complex open-domain questions through iterative query generation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 2590–2602. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Siva Reddy, Oscar Täckström, Michael Collins, Tom Kwiatkowski, Dipanjan Das, Mark Steedman, and Mirella Lapata. 2016. Transforming dependency structures to logical forms for semantic parsing. In *Association for Computational Linguistics (ACL)*.
- Alane Suhr, Stephanie Zhou, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. A corpus for reasoning about natural language grounded in photographs. *Association for Computational Linguistics (ACL)*.
- Alon Talmor and Jonathan Berant. 2018. The web as knowledge-base for answering complex questions. In *North American Association for Computational Linguistics (NAACL)*.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. Constructing datasets for multi-hop reading comprehension across documents. *Transactions of the Association for Computational Linguistics*, 6:287–302.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4:401–415.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan R. Salakhutdinov, and Christopher D. Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Wen-tau Yih, Matthew Richardson, Christopher Meek, Ming-Wei Chang, and Jina Suh. 2016. The value of semantic parse labeling for knowledge base question answering. In *Association for Computational Linguistics (ACL)*.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir R. Radev. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- John M. Zelle and Raymond J. Mooney. 1996. Learning to parse database queries using inductive logic programming. In *Association for the Advancement of Artificial Intelligence (AAAI)*, pages 1050–1055.
- Luke Zettlemoyer and Michael Collins. 2005. Learning to map sentences to logical form: Structured classification with probabilistic categorical grammars. In *Uncertainty in Artificial Intelligence (UAI)*, pages 658–666.