

ASQ: Automatically Generating Question-Answer Pairs Using AMRs

Geetanjali Rakshit and Jeffrey Flanigan
Computer Science and Engineering Department
UC Santa Cruz
`{grakshit, jmflanig}@ucsc.edu`

Abstract

We introduce ASQ, a tool to automatically mine questions and answers from a sentence using the Abstract Meaning Representation (AMR). Previous work has used question-answer pairs to specify the predicate-argument structure of a sentence using natural language, which does not require linguistic expertise or training, and created datasets such as QA-SRL and QAMR, for which the question-answer pair annotations were crowdsourced. Our goal is to build a tool (ASQ) that maps from the traditional meaning representation AMR to a question-answer meaning representation (QMR). This enables construction of QMR datasets automatically in various domains using existing high-quality AMR parsers, and provides an automatic mapping AMR to QMR for ease of understanding by non-experts. A qualitative evaluation of the output generated by ASQ from the AMR 2.0 data shows that the question-answer pairs are natural and valid, and demonstrate good coverage of the content. We run ASQ on the sentences from the QAMR dataset, to observe that the semantic roles in QAMR are also captured by ASQ.

1 Introduction

In the English language, predicate-argument structure of a sentence plays a key role in understanding its meaning. The idea of annotating predicate-argument structure been explored in the past in the creation of linguistic resources such as PropBank (Palmer et al., 2005), FrameNet (Baker et al., 1998), OntoNotes (Pradhan et al., 2007), NomBank (Meyers et al., 2004), Groningen Meaning Bank (GMB) (Bos et al., 2017), *inter alia*. Relational semantics at the sentence level has been encoded in the Abstract Meaning Representation (AMR) (Banarescu et al., 2013) using PENMAN notation. In contrast, He et al. (2015) introduced the idea of using question-answer pairs to repre-

Sentence: Something was broken in my engine .
AMR:
(b / break-01
:ARG1 (s / something
:location (e / engine
:poss (i / i)))
QA pairs:
What was broken? something
Where was something broken? Engine
Whose engine? my

Figure 1: Example from ASQ

sent the semantic roles in a sentence, where each predicate-argument relationship is expressed as a question-answer pair. This idea was further explored in large-scale QA-SRL parsing (FitzGerald et al., 2018), and QAMR (Michael et al., 2017).

Question-answer meaning representations (QMRs) are appealing because they are intuitive and understandable to non-experts. They have been shown to closely match the structure of traditional predicate-argument structure annotation schemes (PropBank, FrameNet, NomBank), without the need for predefined frames or thematic role ontologies. However, QMRs contain the complexities and ambiguities of natural language. On the other hand, standard MRs are unambiguous, and are widely used in downstream tasks, but are not understandable to non-experts.

We would like to be able to map automatically between standard MRs and QMRs, to be able to present the same information either in a format understandable to computers, or in a format understandable to non-expert humans. We take a first step in this direction and demonstrate this is possible by creating an automatic system to map AMR to a QMR. We introduce AMR Sourced Ques-

tions (ASQ), a tool to automatically mine questions and their answers from the AMR for a sentence. Figure 1 shows an example of representing an AMR as question-answer (QA) pairs. We hand-craft templates to convert the AMR into compact, non-redundant QA pairs, with good coverage. We run ASQ on the sentences from the QAMR dataset, to show that the semantic roles in QAMR are also captured by ASQ.

We show that outputs from ASQ have good coverage with minimal redundancy, while being fast and requiring minimal supervision, to produce good quality QA pairs. The strength of ASQ comes from the quality of templates and that of the AMRs.

The main contributions of this work are:

- ASQ: a tool to automatically obtain semantic information in the form of question-answer (QA) pairs, given a sentence and its AMR. To the best of our knowledge, generating question-answer pairs using AMR has not been done before.
- A collection of rich templates for question generation, based on AMR relations.
- A question answering dataset built from AMR 2.0, containing the range of semantics expressed by AMRs.
- Additional question answering annotations, using ASQ and an automatic AMR parser, on sentences from the QAMR dataset.

Our code and data can be found at <https://github.com/geetanjali-rakshit/ASQ>.

2 Related Work

He et al. (2015) introduced the QA-SRL task to specify the predicate argument structure of a sentence using natural language in the form of question-answer pairs. The QA-SRL dataset was expanded to a large scale in (FitzGerald et al., 2018). (Michael et al., 2017) introduced QAMR, which differs from QA-SRL in focusing on all words in the sentence (and not just verbs). These resources are shown to capture the same kinds of predicate-argument structures as PropBank (Palmer et al., 2005), FrameNet (Baker et al., 1998), etc. Meaning representations prior to these include UDS (White et al., 2016), UCCA (Abend and Rapoport, 2013), GMB (Basile et al., 2012), etc.

While both QAMR and QA-SRL are manually created, QA-SRL defines a limited grammar and

templates to help annotators frame questions, and fast track the annotation process. Our work is also highly template-based, but we use richer and more number of templates, and we do not define a grammar for it. QAMR, on the other hand, doesn’t use templates and allows free-form questions instead.

3 Background

3.1 Abstract Meaning Representation

Abstract Meaning Representation is a semantic representation of a whole sentence to convey its logical meaning (Banarescu et al., 2013). It draws heavily from PropBank (Palmer et al., 2005), while adding more semantic information of its own. AMRs do not specify information like tense and aspect.

AMRs are represented as rooted, labeled graphs, which may be written in the graph format, or as a string in the PENMAN notation, or, as logical triples. Figure 1 shows a sentence and its AMR. Each relation in the AMR graph has a **relation name**, a **concept**, and a **variable**. For example, in figure 1, the variables are: *b*, *s*, *e*, *i*; the corresponding *concepts* are: *break-01*, *something*, *engine*, *i*; and the *relations*: *ARG1*, *location* and *pos*. The root doesn’t have a relation name. The numeral *-01* in *break-01* denotes the sense of the verb *break* from PropBank.

Sentences with the same meaning are expected to have the same AMR (although there isn’t always a single unique AMR for a sentence). One AMR can therefore have multiple surface realizations, all meaning the same. The relations express how words in the sentence are related. These include **core roles**, i.e., frame arguments (following PropBank) such as *arg0*, *:arg1*, *:arg2*, *:arg3*, *:arg4*, *:arg5*; and **non-core roles** such as general semantic relations like *:accompanier*, *:age*, *:beneficiary*, *:cause*; relations for quantities like *:quant*, *:unit*, *:scale*; relations for date-entities *:day*, *:month*, *:year*, *:weekday*, *:time*; and relations for lists *:op1*, *:op2*, *:op3*, *:op4*. There are about 100 relations listed in the AMR specifications. These relations are the key to generating question-answer pairs in our approach, as described in the subsequent sections.

4 Approach

Our approach uses the sentence text and its AMR to produce question-answer pairs. We create a resource of handcrafted templates to generate questions from, described in detail in §4.1. We have

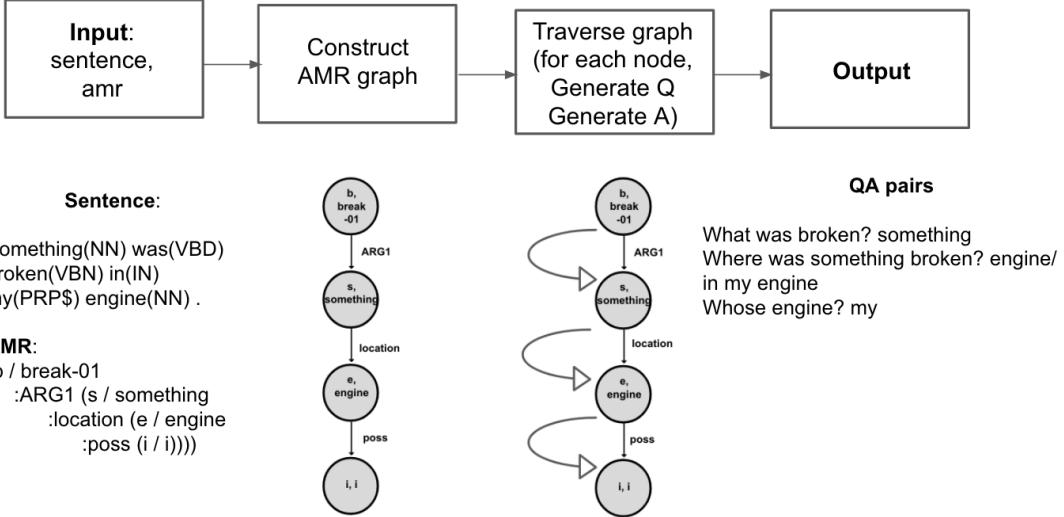


Figure 2: System Overview

at least one template for each relation. These templates are transformed into a suitable question to be answered from the sentence under consideration.

The overall system pipeline is shown in figure 2. We read the AMR as a string in the PENMAN format, and construct an AMR graph over it (§4.2), where each node is labeled with a *variable* name and *concept*, and the edges are labeled with *relations*.

4.1 Template Creation

For the relations in the AMR specifications (Banarescu et al., 2012), we handcraft one or more templates. These templates are then transformed to construct a suitable question that can be answered by the sentence. The distinction between the two kinds of AMR relations, core and non-core, is based on whether or not the predicate information is essential to understand the meaning of the relation of the word to its parent word, or vice versa. Therefore, the construction of templates for each kind of role is handled differently, as described next. Table 1 shows some statistics about the templates created.

The core roles in AMR come from the numbered argument roles in PropBank (Palmer et al., 2005). The meaning of the arguments depend on the predicate, and could have completely different meanings for the same core relation but different predicates. The predicates themselves carry sense annotations from PropBank.

Designing templates for all the predicate-argument combinations is a task in itself. As

Core roles

Number of templates	483
Avg number of templates for each relation	7.3
Non-core roles	
Number of templates	385
Avg number of templates for each relation	2.5

Table 1: Statistics about Templates

a workaround, we use the mappings in Semlink (Palmer, 2009) to associate each numbered argument with a **thematic role** in VerbNet (Kipper et al., 2000). We then create templates for each thematic role in VerbNet. When we have to choose templates based on a numbered argument and predicate, we can find the thematic role that the argument maps to, and use the templates for that role. When creating these templates, we also store the tense of the verb required to convert the template into a valid question. This helps to narrow down on the number of choices when selecting a template to construct questions. Table 2 shows templates for some of the core-roles.

Non-core Roles For the non core roles, the *relation* name is sufficient to characterize templates for it. We do not use predicate information to create these templates. Table 3 shows some of the templates for the core-roles.

Some templates have multiple blanks to be filled. We store the accepted parts-of-speech tag informa-

Rel	Tense	Template
ARG0	past, pres.	Who _____ ?
	past, pres.	What _____ ?
	pres.	Who is _____ ?
	pres.	What is _____ ?
	pres.	Who are _____ ?
	pres.	What are _____ ?
	past	Who was _____ ?
	past	What was _____ ?
	past	Who were _____ ?
	past	What were _____ ?
future	future	What will _____ ?
	future	Who will _____ ?
ARG1	past	Who does someone _____ ?
	past	What does someone _____ ?
	pres.	What is someone _____ ?
	future	What will someone be _____ ?
ARG2	pres.	What is something _____ from ?
	pres.	What is something _____ of ?
ARG3	pres.	Who is something _____ for ?
	pres.	What is something _____ for ?

Table 2: Some templates for the core roles for the predicate *make*. ARG0 is the agent, ARG1 is the product, ARG2 is the material, ARG3 is the beneficiary. Tense is the tense of the generated surface form of *make*.

tion for each blank.

4.2 AMR Graph Preprocessing

We construct a tree over the AMR graph by using the tree structure defined by the PENMAN notation of the annotation. Once this tree is constructed, we perform some steps of preprocessing to avoid redundancy and keep cohesive bits of information together by condensing the children nodes into the parent. For example, when the concept for a node is *temporal-quantity*, we replace it with the concepts from children nodes with relations *:quant* and *:unit* to get a concept like "1 year". We do this for concepts like *date-entity*, *distance-entity*, *area-entity*, *volume-entity*. For leaf nodes with the relation *:op*, we combine the concepts from children at the same level with relation *:op{x}* where x can be 1, 2, 3, etc, which helps to extract multi word proper nouns as a single concept. We ignore relations like *:polarity*, *:wiki*, *,:polite*, *,:polite-of*, *,:mode* for question generation since they often do not contain significant information to pose questions about.

We do a depth-first pre-order traversal of the AMR graph, to process each node to create

question-answer pairs.

4.3 Question Generation

We create a question for every node in the final AMR graph, except for the root node. Figure 3 shows the process of obtaining a question from each node in the AMR graph.

For each node, we check its relation with its parent, i.e., the label on its link to the parent node, and retrieve the templates for that relation from the previously created templates. For core roles, we also use the parent verb information to select the right templates. The next step is to transform the template(s) into suitable questions, which involves substituting the blank(s) with the correct form of the *concept* from the parent node. This gives us a question pertaining to each template. We score these questions using the GPT-2 language model (Radford et al., 2019), and output the question with the highest score.

We illustrate the question generation process with an example in figure 3. From the input (the sentence, and the node and its parent from its AMR graph), we extract the relation *ARG1*. Next, we extract the templates for *ARG1* : *What was _____ ?* and, *What _____ ?* Then, we transform these templates to pose a question for the relation *ARG1*. This requires choosing the right word, i.e., the correct form of the parent concept (in this case *break-01*) that goes into the blank. We lemmatize the sentence, and get the lemma of each word, to determine which lemma corresponds to the parent concept. "broken" is lemmatized to 'break'. Therefore, we substitute the word "broken" from the sentence into the blank. The two templates are thus transformed to two questions: *What was broken ?*, and *What broken ?*. Of these, *What was broken ?* is chosen as the output because the language model gives it a higher score. For each verb in the sentence, we also add a question to capture the sense of the verb, which can be answered by the sense annotation.

4.4 Answer Generation

The answers to the questions are spans of words from the sentence. The answer could be a single word, or a phrase, which comes from the *concept* of the node for which we are trying to generate a question-answer pair. We lemmatize the sentence to obtain the mapping of concepts in the AMR to words in the sentence. Some questions can be answered by the word to which the *concept* lemma

Relation	POS tag	Tense	Template
location	verb	present	Where is _____ ?
	verb	past	Where was _____ ?
	verb	past	Where did someone _____ ?
mod	verb	past	Which _____ ?
	verb	past	What type of _____ ?
frequency	verb, noun	present	How many times someone _____ _____ ?
	verb, noun	present	How many times something _____ _____ ?

Table 3: Templates for Non-core Roles

Question Generation

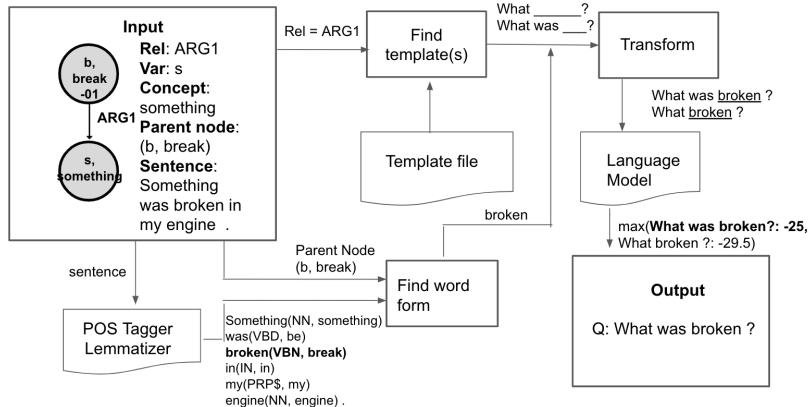


Figure 3: Question Generation

is mapped. However, to extract the correct phrase is not as straightforward. Therefore, we use the Stanza library from (Qi et al., 2020) to get the dependency parse tree of the sentence and then extract the sub-tree rooted at the word to which the *concept* lemma is mapped as the answer. Figure 5 shows the answer generation module.

In figure 4, we show an example where we extract an entire phrase to answer a question. This explains how to obtain the answer for a question posed for the node (*d*, *desert*). The *concept* *desert* maps to the word *desert*. Therefore, in the dependency parse tree, we look for the head *desert*, and extract all the words modified by *desert* and its children, i.e., all the words from the sub-tree, as the answer.

5 Experiments and Results

We run ASQ on the AMR annotated Little Prince Corpus (1,562 sentences) and the AMR Annotation Release 2.0 Corpus (39,260 sentences) (Knight et al., 2014) to get question-answer pairs. Table 4 show some statistics on the output obtained from both of these datasets. Since they contain gold-standard annotations of AMRs, we perform a qual-

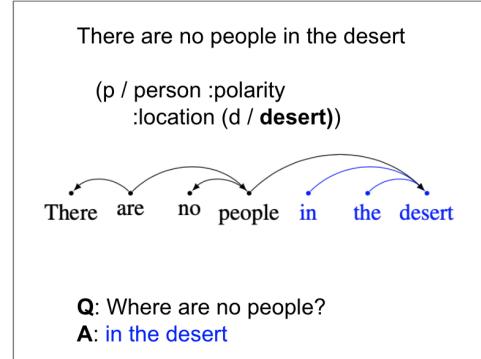


Figure 4: Answer extraction from dependency tree

itative evaluation of the results from AMR 2.0.

5.1 Human Evaluation

We randomly select 70 sentences from the test split of AMR 2.0 and use this as our test set for error analysis.

Qualitative Evaluation We show the sentences from the test set and the question-answer pairs generated from them by ASQ to a test subject and ask them to score the quality of each of the QA pairs for grammatical and naturalness on a scale of

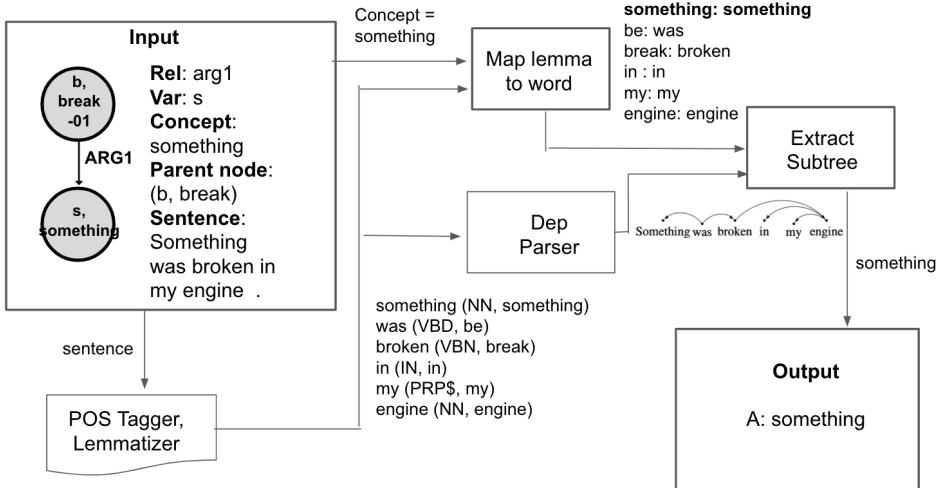


Figure 5: The answer generation pipeline.

Little Prince AMR 2.0		
Total # of questions	23,433	468,263
Avg # of questions per sentence	15.00	11.92
# of unique words	12,496	183,334
Avg length of questions	4.23	3.92
Avg length of answers	3.7	2.2

Table 4: Statistics for the Little Prince Corpus, and the AMR 2.0 corpus.

Coverage	69%
Grammaticality	4.1
Naturalness	3.7
Percentage of correct answers	67.8%

Table 5: Qualitative Evaluation on a test set of 70 random sentences from the test split of AMR 2.0 showing averaged scores.

After that they are not able to move , and they sleep through the six months that they need for digestion .	
Q: What is the meaning of possible in this sentence ?	Q: What sleep ?
A: likely or able to be/occur,	A: they
Q: What is the meaning of move in this sentence ?	Q: What duration ?
A: change location	A: 6 month
Q: What moves ?	Q: What is the meaning of need in this sentence ?
A: they	A: need, required, necessary,
Q: When someone moves ?	Q: What is the purpose of need ?
A: after that	A: digest
Q: What is the meaning of sleep in this sentence ?	Q: What is the meaning of digest in this sentence ?
A: sleep, slumbering,	A: take in completely, like food, breaking down/processing of food

Figure 6: Text from the Little Prince Corpus, with output QA pairs from ASQ

1-5, (5 being the highest), and the correctness of the answer for the question (0 for incorrect, 0.5 somewhat correct, 1 for correct). Also for each sentence, we ask them to rate coverage, i.e., how much information content from the sentence is captured in the set of question-answer pairs generated for it, expressed as a percentage. For the outputs from ASQ, the number of questions is also a good estimate of coverage, since almost all the QA pairs are meant to convey different pieces of information. These results are summarized in Table 5.

Error Propagation from AMR parsers We use the AMR-GS parser (Cai and Lam, 2020) trained on AMR 2.0, which has close to 80% accuracy on the AMR2.0 test set. The sentences in QAMR tend to be longer, and more informative, and are from a different domain. The accuracy of the AMR predictions starts to drop as the sentence length increases. For sentences shorter than 12 tokens, we obtain very similar results for gold standard and predicted AMRs. Figure 7 shows an example of QA pairs from gold standard and predicted AMRs.

Akayev highly commended the results of the conference with Peng Li	
Gold-standard AMR	<p>question: What is the sense of “command” in this sentence ? answer: give praise.</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: What is the sense of “result” in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: To what degree did something commend ? answer: high</p>
\Li_Peng\\	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
\Peng\\:op2 (\Li)))	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
\degree (h / high))	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>

(a) QA Pairs from ASQ constructed from gold-labeled AMRs

Predicted AMR	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
c0 / command-01	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:ARG1 (r0 / result-01	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:ARG1 (c1 / conference))	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:ARG0 (p0 / person	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:name (n0 / name	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:op1 \Peng)	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:op2 (\Li))	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:ARG1-1-of (p1 / person	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:name (n1 / name	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:op1 \Akayev))	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>
:degree (h0 / high-02))	<p>question: What is the sense of “command” in this sentence ? answer: give praise</p> <p>question: What is the sense of \result in this sentence ? answer: come about, become, come about from a prior condition</p> <p>question: What commends ? answer: person</p> <p>question: What was the name of person ? answer: Peng Li</p> <p>question: What was the name of person ? answer: Akayev</p> <p>question: To what degree did something commend ? answer: high</p>

(b) QA Pairs from ASQ constructed from AMRs predicted by AMR-GS parser

Figure 7: For relatively smaller sentences, the output QA pair are similar

5.2 Comparison with QAMR

The QAMR dataset (Michael et al., 2017) contains 100,000 questions from over 5000 sentences, from Wikipedia, WikiNews and the Penn Treebank. These questions were written by annotators manually. In the annotation set up, workers were shown an English sentence with target words highlighted, and they were asked to write *valid* question-answer pairs for each target word. Workers were incentivized to write more questions, to get better coverage.

We use the same definition of a *valid* question as (Michael et al., 2017), which says a question is *valid* if: it contain at least one word from the sentence; is about the sentence’s meaning; can be answered obviously and explicitly in the sentence; not be a yes/no question; and not be redundant. Redundancy is defined informally as two questions having the same meaning and the same answer.

Using AMRs to pose questions automatically ensures that the requirements for a question to

	QAMR	ASQ-MR
Total # of questions	89,091	56,355
Avg # of questions per sentence	18.1	12.15
# of unique words	721,122	225,420
Avg length of questions	16	7.04
Avg length of answers	6.7	3.11

Table 6: Statistics comparing QAMR and automatic ASQ-MR on the QAMR dataset

be *valid* are met. The questions generated are about the meaning of the sentence, and are answerable from information from the sentence. In using AMRs, we automatically ensure maximum coverage, because a good AMR captures most of the factual information in the text. The problem of mapping question-answer pairs to predicate-argument roles becomes redundant, since we create a question-answer pair for almost every relation. Also, there is no scope for redundancy in creating questions from each node of the AMR.

We obtain AMRs for the sentences in the QAMR corpus using an AMR parser (Cai and Lam, 2020). We run ASQ on the sentences from QAMR and their AMRs, to automatically get question-answer pairs, referred to as ASQ-MR. Table 6 lists some statistics for the manually created QAMR and automatic ASQ-MR question-answer pairs. The QAMR annotations have higher number of questions per sentence, however, even with filtering, a lot of questions are redundant. Note that the quality of the QA pairs obtained on the QAMR sentences is observed to be much worse than those obtained on the AMR 2.0 or the Little Prince datasets (discussed further in §5.1).

Table 7 shows examples of semantic roles from QAMR and ASQ-MR. One semantic phenomena that the QAMR data exhibits but ASQ-MR doesn’t is the use of synonyms to substitute words from the text in creating the questions. It is natural for humans to express information in words that they are more familiar with. This results in more diverse questions in QAMR and a richer vocabulary.

6 Conclusion and Future Work

We have presented ASQ, a system for automatically mining questions and their answers from the AMR of a sentence. This enables construction of QMR datasets automatically in various domains using existing high quality AMR parsers, and provides

Semantic role-label	Examples
core relations for open-class concepts	His family’s high rank enabled Ibn Khaldun to study with the best teachers in Maghreb. question: Who ranks ? answer: His family ’s
modifiers	ASEAN members, together with the group’s six major trading partners (Australia, China, India, Japan, New Zealand, South Korea), began the first round of negotiations on 26-28 February 2013, in Bali, Indonesia on the establishment of the Regional Comprehensive Economic Partnership. question: What type of partner ? answer: trade
temporal	Farage, who has lead the party since 2010, announced his resignation on Friday, with the recommendation that Suzanne Evans serve as interim leader. question: When someone announces ? answer: weekday Friday
location	A petroleum company based in India, Oil and Natural Gas Corporation (ONGC), operates in a number of oil blocks under South China Sea with consent of Vietnam, in the same Phu Khanh basin. question: Where is someone operating ? answer: in the same Phu Khanh basin
causality	In the mean time, much stronger friction over land quickly reduces their strength. question Why something reduced ? answer: stronger friction land
negation	His friends and family gained important positions without qualifications. question: How did someone gain? answer: not qualify
quantifiers	People with schizophrenia will sometimes report that, although they are acting in the world, they did not initiate, or will, the particular actions they performed. question: How often report ? answer: sometimes
degree	The Serbian inventor and electrical engineer Nikola Tesla was heavily influenced by Goethe’s Faust, his favorite poem, and had actually memorized the entire text. question: To what degree did someone influence ? answer: heavy
discourse information	KH: John and Russ (Nash in the film) went around [scouting] but to be honest, Russ’s dad owned the land where we filmed 90% of the movie. question: What contrast? answer: KH John went around scouting owned .

Table 7: Examples of semantic roles in ASQ-MR that are also expressed in QAMR

an automatic mapping AMR to QMR for ease of understanding by non-experts.

For future work, it would be possible to improve ASQ by the addition of more templates. ASQ could also be extended to include the use of synonyms. For the data from ASQ, post-editing the questions make an improved version of the dataset. This would require little effort, as it would be easier and quicker for annotators to fix a somewhat incorrect question, than to have to come up with questions by themselves.

Limitations

While ASQ demonstrates the feasibility of automatically generating question–answer pairs from AMRs, several limitations remain. ASQ’s performance is tightly coupled with the accuracy of the input AMRs. When gold-standard AMRs are available, the system produces high-coverage and natu-

ral questions. However, with automatically parsed AMRs—particularly for longer or domain-shifted sentences (e.g., QAMR)—errors in parsing propagate directly to question generation, leading to reduced grammaticality, incorrect answers, and occasional semantic drift. Although ASQ uses an extensive handcrafted template library for AMR relations, these templates cannot fully capture the diversity and nuance of natural question phrasing. The reliance on fixed patterns can result in repetitive structures and limited syntactic variation. Furthermore, certain relations or semantic phenomena may be underrepresented in the template set, affecting completeness. Unlike human-authored datasets such as QAMR, ASQ does not currently incorporate paraphrasing or synonym substitution. This limits the lexical diversity of generated questions and may reduce their naturalness or alignment with user expectations in real-world applications.

References

Omri Abend and Ari Rappoport. 2013. Universal conceptual cognitive annotation (ucca). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 228–238.

Collin F Baker, Charles J Fillmore, and John B Lowe. 1998. The berkeley framenet project. In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 86–90.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2012. Abstract meaning representation (amr) 1.0 specification. In *Parsing on Freebase from Question-Answer Pairs. In Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing. Seattle: ACL*, pages 1533–1544.

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 *Proceedings of the 7th linguistic annotation workshop and interoperability with discourse*, pages 178–186.

Valerio Basile, Johan Bos, Kilian Evang, and Noortje Venhuizen. 2012. Developing a large semantically annotated corpus. In *Eighth International Conference on Language Resources and Evaluation*, pages 3196–3200. EUROPEAN LANGUAGE RESOURCES ASSOC-ELRA.

Johan Bos, Valerio Basile, Kilian Evang, Noortje J Venhuizen, and Johannes Bjerva. 2017. The groningen meaning bank. In *Handbook of linguistic annotation*, pages 463–496. Springer.

Deng Cai and Wai Lam. 2020. Amr parsing via graph-sequence iterative inference. *arXiv preprint arXiv:2004.05572*.

Nicholas FitzGerald, Julian Michael, Luheng He, and Luke Zettlemoyer. 2018. Large-scale qa-srl parsing. *arXiv preprint arXiv:1805.05377*.

Luheng He, Mike Lewis, and Luke Zettlemoyer. 2015. *Question-answer driven semantic role labeling: Using natural language to annotate natural language.* In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 643–653, Lisbon, Portugal. Association for Computational Linguistics.

Karin Kipper, Hoa Trang Dang, Martha Palmer, et al. 2000. Class-based construction of a verb lexicon. *AAAI/IAAI*, 691:696.

Kevin Knight, Lauren Baranescu, Claire Bonial, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Daniel Marcu, Martha Palmer, and Nathan Schneider. 2014. Abstract meaning representation (amr) annotation release 1.0. *Web download*.

Adam Meyers, Ruth Reeves, Catherine Macleod, Rachel Szekely, Veronika Zielinska, Brian Young, and Ralph Grishman. 2004. The nombank project: An interim report. In *Proceedings of the workshop frontiers in corpus annotation at hlt-naacl 2004*, pages 24–31.

Julian Michael, Gabriel Stanovsky, Luheng He, Ido Dagan, and Luke Zettlemoyer. 2017. Crowdsourcing question-answer meaning representations. *arXiv preprint arXiv:1711.05885*.

Martha Palmer. 2009. Semlink: Linking propbank, verbnet and framenet. In *Proceedings of the generative lexicon conference*, pages 9–15. GenLex-09, Pisa, Italy.

Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational linguistics*, 31(1):71–106.

Sameer S Pradhan, Eduard Hovy, Mitch Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2007. Ontonotes: A unified relational semantic representation. In *International Conference on Semantic Computing (ICSC 2007)*, pages 517–526. IEEE.

Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. Stanza: A “python” natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.

Aaron Steven White, Drew Reisinger, Keisuke Sakaguchi, Tim Vieira, Sheng Zhang, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2016. Universal decompositional semantics on universal dependencies. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1713–1723.