

# Addressing Community Question Answering in English and Arabic

Giovanni Da San Martino, Alberto Barrón-Cedeño, Salvatore Romeo, Alessandro Moschitti, Shafiq Joty, Fahad A. Al Obaidli  
Qatar Computing Research Institute  
HBKU, Doha, Qatar  
{albarron,gmartino,sjoty, faalobaidli, amoschitti,sromeo}@qf.org.qa

Kateryna Tymoshenko, Antonio Uva, Daniele Bonadiman  
Department of Computer Science and Information Engineering,  
University of Trento  
Trento, Italy  
{kateryna.tymoshenko, antonio.uva, d.bonadiman}@unitn.it

## ABSTRACT

This paper studies the impact of different types of features applied to learning to re-rank questions in community Question Answering. We tested our models on two datasets released in SemEval-2016 Task 3 on “Community Question Answering”. Task 3 targeted real-life Web fora both in English and Arabic. Our models include bag-of-words features (BoW), syntactic tree kernels (TKs), rank features, embeddings, and machine translation evaluation features. To the best of our knowledge, structural kernels have barely been applied to the question reranking task, where they have to model paraphrase relations. In the case of the English question re-ranking task, we compare our learning to rank (L2R) algorithms against a strong baseline given by the Google-generated ranking (GR). The results show that (i) the shallow structures used in our TKs are robust enough to noisy data and (ii) improving GR is possible, but effective BoW features and TKs along with an accurate model of GR features in the used L2R algorithm are required. In the case of the Arabic question re-ranking task, for the first time we applied tree kernels on syntactic trees of Arabic sentences. Our approaches to both tasks obtained the second best results on SemEval-2016 subtasks B on English and D on Arabic.

## 1. INTRODUCTION

In recent years, there has been a renewed interest in information retrieval for community question answering (cQA). This combines traditional question answering with a modern Web scenario, where users pose questions on a forum expecting to receive answers from other users. As fora are fairly open, many contributors may post comments loosely connected to the original question, whereas others are simply wrong or are not answers at all. We addressed this problem in [3, 13, 20].

One of the most critical problem arises when a user posts a new question to the forum. In this case, the retrieval system would search for relevant comments linked to other questions in order to find potentially appropriate answers. In this noisy and complex set-

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

SIGIR '16, July 17 - 21, 2016, Pisa, Italy

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4069-4/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2911451.2914749>

ting, also powerful search engines (e.g., Google), have hard time to retrieve comments that can correctly answer the original question. An approach to deal with this case is to divide the problem into two separate tasks: (i) retrieving a set of similar questions and their associated set of comments and (ii) assessing the usefulness of the retrieved comments with respect to the question posed by the user. In this paper we focus on first task both in English and Arabic. This shows that our models can be applied to different languages.

The retrieval of relevant questions requires models different from typical search engines. Firstly, we deal with short texts: questions may include descriptions or subquestions but they are not typically larger than one or two paragraphs. Thus rich similarity features are necessary to deal effectively with the sparseness of such texts. Secondly, the syntactic structure of the questions is rather important for learning to recognize paraphrases and thus selecting the right candidates. For instance, consider the following two questions :

$q_1$  How do I get a visa for Qatar to visit my wife?

$q_2$  How do I get a visa for my wife to have her visit Qatar?

Question  $q_1$  has roughly the same BoW representation as question  $q_2$ , but their information request—therefore their answers—are totally different. The structure of the questions can be exploited to detect the difference between these difficult cases.

In the past, the study of cQA models has been limited by a substantial lack of manually annotated data for question-question similarity, as the main annotation was provided by the user themselves and thus essentially not fully reliable. Recently, a new resource has been released for the SemEval 2016 Task 3 on Answer Selection in cQA [19].<sup>1</sup> Given a set of existing forum questions  $Q$ , where each existing question  $q \in Q$  is associated with a set of answers  $C_q$ , and a new user question  $q'$ , the ultimate task is to determine whether a comment  $c \in C_q$  represents a pertinent answer to  $q'$  or not. This task can be subdivided into three tasks, namely: (A) to assign a relevance (*goodness*) score to each answer  $c \in C_q$  with respect to the existing question  $q$ ; (B) to re-rank the set of questions  $Q$  according to their relevance against the new question  $q'$ ; and finally (C) to predict the appropriateness of the answers  $c \in C_q$  against  $q'$ . An adaptation of Task C was proposed for Arabic (Task D). In this paper, we focus on tasks B and D: question re-ranking in both English and Arabic.

For task B we use (i) text similarity features, derived from the classical BoW representations, e.g., n-grams, skip-grams; (ii) syntactic/structural features injected by TKs, which have shown to

<sup>1</sup><http://alt.qcri.org/semeval2016/task3>

achieve the state of the art in the related task of answer sentence retrieval [27]; and (iii) features for modeling the initial rank provided by a state-of-the-art search engine, which represents a strong baseline. Our extensive experimentation, produced the following results: (i) the BoW features based on similarity measures alone do not improve GR. (ii) Our TKs applied to questions alone outperform the models based on similarity measures and when jointly used with the rank features improve all the models. In particular, they outperform GR by 1.72 MAP points (95% of statistical confidence).

For task D we apply tree kernels on pairs of syntactic trees of Arabic sentences and define additional features derived from machine translation evaluation scores.

Our approaches to both tasks B and D granted us the second position on SemEval 2016 Task 3.

The rest of our contribution is distributed as follows. Section 2 includes related work. Section 3 describes the approached problem. Section 4 describes our learning to rank models. Section 5 discusses our experiments and obtained results. Section 6 includes conclusions and further work.

## 2. RELATED WORK

The first step for any system that aims to automatically answer questions on cQA sites is to retrieve a set of questions similar to the user’s input. The set of similar questions is later used to extract possible answers for the input question. However, determining question similarity remains one of the main challenges in cQA due to problems such as the "lexical gap". To overcome this problem, different approaches have been proposed. Early methods used statistical machine translation techniques to compute semantic similarity between two questions. For instance, [31] applied a phrase-based translation model. Their experiments on Yahoo! Answers showed that models based on phrases are more effective than those using words, as they are able to capture contextual information. However, approaches based on SMT have the problem of requiring lots of data in order to estimate parameters.

Algorithms that try to go beyond simple text representation are presented in [5] and [7]. In [5] a similarity between two questions on Yahoo! Answers is computed by using a language model with a smoothing method based on the category structure of Yahoo! Answers. In [7], the authors search for semantically similar questions by identifying the topic and focus of the user’s question. More specifically, they compute a similarity between the questions’ topic, which represents general users interests, and the questions’ focus. A different approach using topic modeling for question retrieval was introduced in [10] and [30]. Here, the authors use LDA topic modeling to learn the latent semantic topics that generate question/answer pairs and use the learned topics distribution to retrieve similar historical questions. The quality of the ranking returned by all these systems was measured on a set of test questions from Yahoo! Answers, with question relevancy judgment annotated by users, sometimes assigned automatically based on heuristics.

It should be noted that the methods above exploited language models or general knowledge given by Yahoo! Answers categories or LDA topics, whereas in our paper, we model the syntactic/semantic relations between pairs of questions using shallow syntactic parsing and lexical matching. The most similar work to ours is [28], where the authors found semantically related questions by computing the similarity between the syntactic trees of the two questions. They used a tree similarity computed in terms of the number of substructures shared between two trees. Different from such approach, we use pairs of questions,  $(q_o, q_s)$ , as learning instances, thus defining

**Table 1: A re-ranking example for the English Question–Question Similarity dataset. For each candidate the Google rank (G), the binary gold standard (GS) relevance, and our rank (R) are reported.**

G	GS	R	Question Text
1	-1	8	The Qatar banana island will be transferred by the end of 2013 to 5 stars resort called Anantara. Has anyone seen this island? Where is it? Is it near to Corniche?
2	+1	2	Is there a good place here where I can spend some quality time with my friends?
3	-1	7	Where is the best beach in Qatar? Maybe a silent and romantic bay? Where to go for it?
4	-1	9	Any suggestions on what are the happenings in Qatar on Holidays? Something new and exciting suggestions please?
5	-1	3	Where in Qatar is the best place for Snorkeling? I’m planning to go out next friday but don’t know where to go.
6	-1	6	Can you give me some nice places to go or fun things to do in Doha for children 17-18 years old? Where can we do some watersports (just for once, not as a member), or some quad driving? Let me know please. Thanks.
7	+1	1	Which all places are there for tourists to Qatar? My nephew 18 years on visit.
8	-1	10	Could you suggest the best holiday destination in the world?
9	-1	5	I really would like to know where the best place to catch fish here in Qatar is. But of course from the beach. I go every week to Umsaeed but rerly i catch somthing! So experianced people your reply will be appreciated.

relational models connecting the syntactic trees of  $q_o$  and  $q_s$ . In this way, the learning algorithms learn transformations that suggest if questions constituted by similar words have similar (paraphrases) or different semantics.

## 3. PROBLEM DESCRIPTION

Among the four task of the SemEval-2016 competition, we focus our attention on the two dealing with question–question similarity. In task B English we consider the questions’ text only. In task D Arabic we also use the information from the answer linked to the forum question. Tasks B uses English instances extracted from *Qatar Living*, a forum for people to pose questions on multiple aspects of daily life in Qatar.<sup>2</sup> Task D uses Arabic instances extracted from three medical fora: *webteb*, *altibbi*, and *consult islamweb*.<sup>3</sup> As these are re-ranking tasks, we evaluate our models using mean average precision (MAP); which is the official evaluation measure of the SemEval 2016 task.

### 3.1 Task B: English Question–Question Similarity

Conceptually, question retrieval is not much different from a standard retrieval task. Given the asked (original) question  $q_o$ , a search engine seeks the Web (or a specific Web forum) for relevant webpages. In cQA, webpages are threads containing questions  $q_s \in Q$ , with their user comments, where the latter can provide information for answering  $q_o$ . For example, Table 1 shows an original

<sup>2</sup><http://www.qatarliving.com/forum>

<sup>3</sup><https://www.webteb.com/>, <http://www.altibbi.com/>, and <http://consult.islamweb.net>.

**Table 2: Class distribution in the training, development, and test partitions for the English Question–Question similarity dataset.**

Class	train	dev	test	overall
Relevant	1,083	214	233	1,530
Irrelevant	1,586	286	467	2,339
Total	2,669	500	700	3,869

**Table 3: Distribution of Relevant and Irrelevant comments at different R ranking positions of GR for the Question–Question similarity dataset.**

R	train	dev	test	overall
1	0.21 ± 0.05	0.24 ± 0.07	0.40 ± 0.11	0.25 ± 0.07
2	0.14 ± 0.03	0.18 ± 0.02	0.12 ± 0.02	0.14 ± 0.03
3	0.11 ± 0.02	0.10 ± 0.01	0.08 ± 0.01	0.10 ± 0.02
4	0.12 ± 0.03	0.08 ± 0.01	0.10 ± 0.03	0.11 ± 0.03
5	0.09 ± 0.02	0.09 ± 0.01	0.08 ± 0.02	0.09 ± 0.02
6	0.08 ± 0.02	0.09 ± 0.02	0.05 ± 0.01	0.08 ± 0.02
7	0.08 ± 0.02	0.07 ± 0.01	0.05 ± 0.01	0.07 ± 0.02
8	0.06 ± 0.01	0.04 ± 0.01	0.03 ± 0.00	0.05 ± 0.01
9	0.07 ± 0.02	0.06 ± 0.01	0.04 ± 0.01	0.07 ± 0.02
10	0.05 ± 0.01	0.05 ± 0.01	0.04 ± 0.01	0.05 ± 0.01

question followed by some questions retrieved by a search engine. In the SemEval dataset we used in our experiments, each question is associated with 10 related threads that are the top 10 webpages retrieved by Google. The main difference with standard document retrieval is the document scoring function. Indeed, although both question and comments are part of the candidate webpage (thread), the question text provides a more synthetic and precise information to infer whether the candidate thread is relevant for  $q_o$ . Table 2 gives class-distribution statistics of the English dataset [19]. A forum question can be a perfect match, relevant, or irrelevant with respect to the new question. For evaluation purposes, both `perfect match` and `relevant` instances are considered `relevant` and must be ranked on top of the `irrelevant` questions. The corpus is composed of 387 user questions, each of which includes 10 potentially related questions. The task organizers used the Google search engine, which represents also the strong baseline for the task, to select potentially relevant forum questions. Table 3 shows the distribution of relevant/irrelevant forum questions per ranking position. Although relevant questions tend to be concentrated towards the top of the Google Rank, they are fairly spread over the entire ranking.

### 3.2 Task D: Arabic Question–Comment Pairs Re-Ranking

A new question and a set of thirty forum question–comment pairs are provided—the comment is always a correct answer to the forum question. The task consists in re-ranking the question–comment pairs according to three classes: (i) `direct`: if it is a direct answer to the new question; (ii) `relevant`: if it is not a direct answer to the question but it provides information related to the topic; and (iii) `irrelevant`: if it is an answer to another question, not related to the topic. For evaluation purposes, both `direct` and `relevant` forum questions are considered as `good`. Table 4 reports statics on the class distribution of the data. There are about 60% of forum questions classified as `good`, thus the dataset is not very unbalanced.

As in the case of English, the Arabic collection questions are user-generated. Since the latter was extracted from medical-domain fora, they show specific challenges: a mix of medical terminology (used by the physicians) and colloquial language (used by the patients). The texts are usually long (on average questions are 50 and

**Table 4: Class distribution in the training, development, and test partitions for the Arabic Question–Comment pairs dataset.**

Class	train	dev	test	overall
direct+relevant	18,329	7634	7619	33582
irrelevant	12,082	5868	5951	23901
Total	30,411	13502	13570	57483

comments are 120 words long). These characteristics pose specific challenges to some of the models which are effective for task B on English, i.e. the tree kernels.

## 4. OUR L2R MODELS

The ranking function for both tasks can be implemented by a scoring function  $r : Q \times Q \rightarrow \mathbb{R}$ , where  $Q$  is the set of questions. The function  $r$  can be linear:  $r(q_o, q_s) = \vec{w} \cdot \phi(q_o, q_s)$ , where  $\vec{w}$  is a linear model and  $\phi()$  provides a feature vector representation of the pair,  $q_o, q_s$ .

We adopt binary SVMs to learn  $r$  from examples and ee model  $\phi(q_o, q_s)$  with different feature sets, which we describe in Sections 4.1 and 4.2.

### 4.1 Tree Kernel Models

Tree kernels are functions that measure the similarity between tree structures. We essentially used the model of [25], originally proposed to rank passages. Different from [25], our questions may contain multiple subquestions, a subject, greetings, and elaborations, thus they are composed of several sentences. We merge the whole question text in a macro-tree using a fake root node connecting the parse trees of all the sentences. In both tasks we represent pairs of questions; therefore, we connect the constituents of two macro-trees corresponding to  $(q_o, q_s)$ , respectively. Figure 1 shows an example where the relations between two Arabic questions are exploited to build a graph. The translation of the two questions is:

$q_o$  “What are the symptoms of irritable bowel syndrome (IBS)?”

$q_s$  “What are the symptoms of irritable bowel?”

We link the two macro-trees by connecting phrases, e.g., NP, VP, PP, when there is at least lexical matching between the phrases of  $q_o$  and  $q_s$ . Note that such links are marked with the presence of a REL tag. Finally, we apply either a partial tree kernel (PTK) or the syntactic tree kernels (STK) [18] and obtain the following kernel:

$$K((q_o, q_s^i), (q_o, q_s^j)) = TK(t(q_o, q_s^i), t(q_o, q_s^j)) + TK(t(q_s^i, q_o), t(q_s^j, q_o))$$

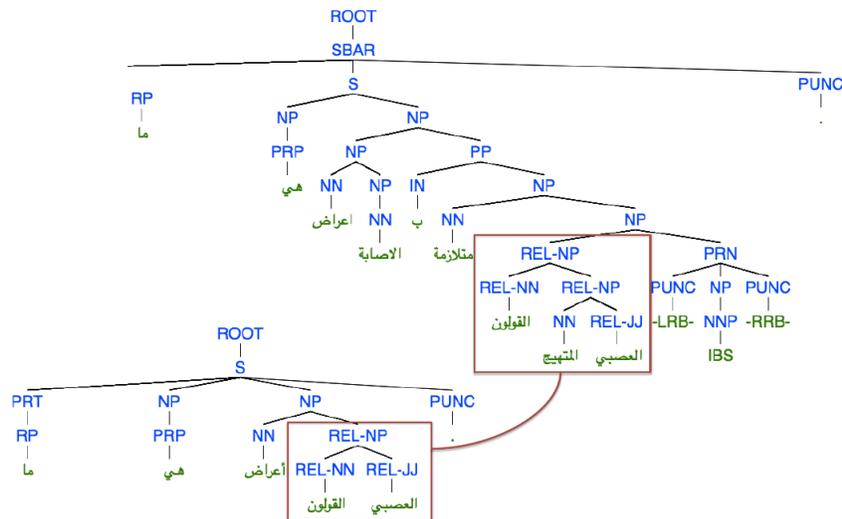
where  $t(x, y)$  extracts the syntactic tree from text  $x$ , enriching it with REL tags computed with respect to  $y$ . Thus  $t$  is an asymmetric function.

For task B we used the PTK kernel. Since the trees of the Arabic data are rather large and very noisy, we used STK, which is faster and uses less features.

### 4.2 Feature Vectors

Our L2R approach relies on various subsets of features to derive the relationship between two texts: text similarities, PTK similarity, Google rank, embedding features and machine translation evaluation features. In task B we used text similarities, PTK similarity, and Google rank. In task D we use text embedding features and machine translation evaluation features.

**Text Similarity Features.** We compute a total of 20 similarities,  $sim(q_o, q_j)$ , using word  $n$ -grams ( $n = [1, \dots, 4]$ ), after stopword



**Figure 1: Our representation based on syntactic trees for the  $q_o$ - $q_s$  pairs enriched with REL links. The translations are  $q_o$  (sentence on top) “What are the symptoms of irritable bowel syndrome (IBS)?” and  $q_s$  “What are the symptoms of irritable bowel?”.**

removal, using greedy string tiling [29], longest common subsequences [1], Jaccard coefficient [9], word containment [15], and cosine similarity.

**PTK Features.** Another similarity is obtained by comparing syntactic trees with PTK, i.e.,  $TK(t(q_o, q_s^i), t(q_s^i, q_o))$ . Note that, different from the model in Section 4.1, PTK here is applied to the members of the same pair and thus only produces one feature.

**Ranking-based Features.** Our ranking feature is based on the ranking generated by Google. Each forum question is located in one position in the range  $[1, \dots, 10]$ . We try to exploit this information in two ways “as-is” ( $pos$ ) or the inverse ( $pos^{-1}$ ).

**Embeddings** We utilize the embedding vectors as obtained by [4]: employing word2vec [17] on the Arabic Gigaword corpus [22]. More specifically, we concatenate the vectors representing a new question and an existing question in the question-answer pair, which is then fed to the SVM classifier.

**MTE features** We used machine translation evaluation features: BLEU [21], TER [26], Meteor [14], NIST [6], Precision and Recall, and length ratio between the question and the comment.

## 5. EXPERIMENTS

We follow the evaluation framework of SemEval 2016 Task 3 in order to be able to compare our system with the ones of the competition [19].

We used different tools for preprocessing the text in English: OpenNLP’s tokenizer, POS-tagger and chunk annotator<sup>4</sup>, and Stanford’s lemmatizer [16], all accessible through DKPro Core [8]<sup>5</sup>. For Arabic texts were first removed stop-words, keeping only content and Latin words. We used the MADAMIRA toolkit [23] for segmenting the texts. In order to split the texts into sentences, we used the Stanford splitter.<sup>6</sup> For parsing Arabic texts into syntactic trees, we used the Berkeley parser [24].

### 5.1 Experiments on Task B

In a set of preliminary experiments, we first compared a reranker

<sup>4</sup><https://opennlp.apache.org/>

<sup>5</sup><https://dkpro.github.io/dkpro-core/>

<sup>6</sup><http://stanfordnlp.github.io/CoreNLP>

—SVMrank [12]— with a standard binary SVM [11]. As the results were comparable, we employed SVMs using the KeLP toolkit<sup>7</sup>, which enables to combine our three subsets of features within different kernels; namely RBF for the similarity features, tree kernels for the parse trees, and either linear or RBF kernels for the ranking-based feature. We set the C parameter of the SVMs to 1 in all the experiments and the parameters of the TKs and RBF kernels to default values.

We conducted three experiments with growing complexity for assessing the effectiveness of our different feature sets (see Section 4), with respect to GR. In agreement, with the SemEval challenge, we evaluate our rankings with Mean Average Precision (MAP), average Recall (AvgRec), and Mean Reciprocal Rank (MRR).

We tested the performance of each of the feature sets in isolation and pair-wise. Table 5 reports the obtained performance both on the development and test sets. The baseline is computed on GR, which produces a strong value as it is a product of the Google technology and its associated knowledge bases. We have two advantages, though: (i) Google is not tuned up on the specific forum data we use and (ii) probably Google does not use syntactic structures in algorithms such as TKs.

Our results support the hypotheses above. Indeed, the MAP of the models derived by Similarities, TK, and their combinations is below GR: without accessing to the Google resources, our models can just approach the search engine’s performance. However, when using the ranking feature, our best model outperforms the MAP of GR by 2.30 and 1.66 absolute percent points on the development and test sets, respectively. The RBF kernel on the ranking feature produces a larger improvement as it can more effectively express higher similarity values when the rankings of questions are close. This cannot be done with a linear kernel.

To better study the result above, Table 6 reports the combinations between the kernel function (Linear or RBF) and the representation of the ranking feature (the ranking itself or its inverse). While all combinations improve above the baseline, there is no clear indication on the choice between  $pos$  or  $pos^{-1}$ . However, the use of the RBF kernel results in the highest performance.

UH-PRHLT-primary, the best system at competition time [19],

<sup>7</sup><https://github.com/SAG-KeLP>

**Table 5: Ranking-based features combined with linear and RBF kernels on the Question–Question similarity dataset.** † shows statistically significant results (at 95%) with respect to GR.

Model	DEV			TEST		
	MAP	AvgRec	MRR	MAP	AvgRec	MRR
GR baseline	71.35†	86.11	76.67†	74.75†	88.30	83.79
Sim.	64.80	82.52	73.73	70.70	85.78	80.58
TK	69.97	86.86	77.73	73.98	88.90	82.55
TK + Sim	71.07	87.72	78.14	73.81	89.21	82.86
<b>Linear Kernel</b>						
Sim + <i>pos</i>	68.04	85.07	76.00	71.99	87.92	81.19
Sim + <i>pos</i> <sup>-1</sup>	70.17	85.98	78.17	75.15	89.19	84.29
TK + <i>pos</i>	71.77	88.46	78.12	75.34	90.67	83.19
TK + <i>pos</i> <sup>-1</sup>	72.64	87.69	75.58	76.18†	90.62	84.62
<b>RBF Kernel</b>						
Sim. + <i>pos</i>	70.42	86.38	78.50	74.61	89.10	83.81
Sim. + <i>pos</i> <sup>-1</sup>	69.82	85.91	77.17	74.58	89.09	83.57
TK + <i>pos</i>	72.93	87.95	77.54	75.72	90.80	83.86
TK + <i>pos</i> <sup>-1</sup>	73.65†	88.78	79.58†	76.41†	91.14	84.62

**Table 6: Performance of different rank features (all models include RBF kernel on similarities and tree kernel) on the Question–Question similarity dataset.** † shows statistically significant results (at 95%) with respect to GR.

Model	DEV			TEST		
	MAP	AvgRec	MRR	MAP	AvgRec	MRR
GR baseline	71.35†	86.11	76.67	74.75†	88.30	83.79
<b>Linear Kernel</b>						
<i>pos</i>	72.18	88.41	78.00	75.67	90.77	83.38
<i>pos</i> <sup>-1</sup>	73.28†	88.47	80.00	76.28†	90.72	84.62
<b>RBF Kernel</b>						
<i>pos</i>	73.24†	88.37	78.40	76.47†	90.78	84.21
<i>pos</i> <sup>-1</sup>	73.60†	88.85	79.67	75.89	90.57	84.14

obtained a MAP of 76.70. This value is not statistically different from our 76.47; still that system uses knowledge bases, such as BabelNet and FrameNet [2], which we do not require.

## 5.2 Experiments on Task D

We have performed three sets of experiments for task D, which coincide with our submissions at the competition and which are reported in Table 7. In the first experiment, *cont*<sub>1</sub>, we applied an SVM with a linear kernel on the embedding features. A second experiment, *cont*<sub>2</sub>, also includes machine translation evaluation features (both features are described in Section 4.2). Finally, a third experiment, *primary*, used a linear kernel on the embedding features in combination with the syntactic tree kernel (STK) [18], applied as described in Section 4.1, to the constituency trees of the question texts.

The submission *cont*<sub>1</sub>, using embedding features from [4], is an average system. When we add the machine translation evaluation (MTE) features the MAP increases from 38.33 to 39.98, making us jump from the 11th to the 7th position in the competition. However, when we combine tree kernels with embedding features MAP improves by more than 7 absolute points (45.50), achieving the second position in the competition, very close to the best system, i.e. 45.83.

The results on the Arabic dataset give more evidence that exploiting syntactical information with tree kernels is crucial for achieving state-of-the-art performance and that tree kernels can be effective for different languages.

## 6. CONCLUSIONS

Task D	MAP	AvgRec	MRR	Rank
primary	45.50	50.13	52.55	2
<i>cont</i> <sub>1</sub>	38.33	42.09	43.75	11
<i>cont</i> <sub>2</sub>	39.98	43.68	46.41	7
best	45.83	51.01	53.66	1
baseline	28.88	28.71	30.93	-

**Table 7: Performance of our primary and contrastive submissions to SemEval-2016 Task D. Best-performing and baseline systems included for comparison. Rank stands for the position in the challenge ranking. Baselines as provided reported in [19].**

We have tackled the task of question re-ranking in English and Arabic using different sets of features and exploiting the syntactical structure of the text via tree kernels. We showed that the combination of similarity features, syntactic structures based on tree kernels and features based on the ranking of search engines is able to boost the performance of a question re-ranker on a real-world cQA dataset. In particular, our results suggest that Google uses general models that can be on par with specific models trained on specific domains. However, if we also use advanced syntactic/semantic representations for modeling the structural relations between questions, we can achieve better results. Moreover, we modeled and tested relational tree kernels for cQA, which are robust to noise and can thus boost Google’s ranking. In the future, we would like to better structure the representation of the questions. Indeed, as mentioned before, there are several different sections of the question text, e.g., subquestions, subject, elaborations. These could be used to improve our shallow representation, which, at the moment, merges all the question trees in a flat macro-tree.

In the comment reranking task we used for the first time tree kernels on an Arabic dataset and showed that they are a key component for reaching a performance comparable to the best SemEval-2016 system. In the future we plan to combine our system with the best one of the competition to further improve our results. Furthermore, we plan to address the challenges caused by the large size of the trees in order to be able to use more powerful tree kernels effectively.

## Acknowledgments

This research is developed by the Arabic Language Technologies (ALT) group at the Qatar Computing Research Institute (QCRI), HBKU, Qatar Foundation in collaboration with MIT. It is part of the Interactive sYstems for Answer Search (IYAS) project.

## 7. REFERENCES

- [1] L. Allison and T. Dix. A bit-string longest-common-subsequence algorithm. *Inf. Process. Lett.*, 23(6):305–310, Dec. 1986.
- [2] C. F. Baker, C. J. Fillmore, and J. B. Lowe. The berkeley framenet project. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics - Volume 1*, ACL ’98, pages 86–90, Stroudsburg, PA, USA, 1998. Association for Computational Linguistics.
- [3] A. Barrón-Cedeño, S. Filice, G. Da San Martino, S. Joty, L. Màrquez, P. Nakov, and A. Moschitti. Thread-level information for comment classification in community question answering. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*

- and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers), pages 687–693, Beijing, China, July 2015. Association for Computational Linguistics.
- [4] Y. Belinkov, M. Mohtarami, S. Cyphers, and J. Glass. VectorSLU: A continuous word vector approach to answer selection in community question answering systems. In *Proceedings of the 9th International Workshop on Semantic Evaluation*, SemEval '15, Denver, Colorado, USA, 2015.
- [5] X. Cao, G. Cong, B. Cui, C. S. Jensen, and C. Zhang. The use of categorization information in language models for question retrieval. In *CIKM*, pages 265–274, 2009.
- [6] G. Doddington. Automatic evaluation of machine translation quality using n-gram co-occurrence statistics. In *Proceedings of the Second International Conference on Human Language Technology Research*, HLT '02, pages 138–145, San Francisco, CA, USA, 2002. Morgan Kaufmann Publishers Inc.
- [7] H. Duan, Y. Cao, C.-Y. Lin, and Y. Yu. Searching questions by identifying question topic and question focus. In *ACL*, pages 156–164, 2008.
- [8] R. Eckart de Castilho and I. Gurevych. A broad-coverage collection of portable nlp components for building shareable analysis pipelines. In *Proceedings of the Workshop on Open Infrastructures and Analysis Frameworks for HLT*, Dublin, Ireland, August 2014.
- [9] P. Jaccard. Étude comparative de la distribution florale dans une portion des Alpes et des Jura. *Bulletin del la Société Vaudoise des Sciences Naturelles*, pages 547–579, 1901.
- [10] Z. Ji, F. Xu, B. Wang, and B. He. Question-answer topic model for question retrieval in community question answering. In *CIKM*, pages 2471–2474, 2012.
- [11] T. Joachims. Making Large-scale Support Vector Machine Learning Practical. In *Advances in Kernel Methods*. MIT Press, Cambridge, MA, USA, 1999.
- [12] T. Joachims. Optimizing search engines using clickthrough data. *KDD*, pages 133–142, 2002.
- [13] S. Joty, A. Barrón-Cedeño, G. Da San Martino, S. Filice, L. Màrquez, A. Moschitti, and P. Nakov. Global thread-level inference for comment classification in community question answering. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 573–578, Lisbon, Portugal, September 2015. Association for Computational Linguistics.
- [14] A. Lavie and M. Denkowski. The METEOR metric for automatic evaluation of machine translation. *Machine Translation*, 23(2–3):105–115, 2009.
- [15] C. Lyon, J. Malcolm, and B. Dickerson. Detecting short passages of similar text in large document collections. *EMNLP*, pages 118–125, 2001.
- [16] C. D. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60, 2014.
- [17] T. Mikolov, W.-t. Yih, and G. Zweig. Linguistic Regularities in Continuous Space Word Representations. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, NAACL-HLT '13, pages 746–751, Atlanta, GA, USA, 2013.
- [18] A. Moschitti. Efficient Convolution Kernels for Dependency and Constituent Syntactic Trees. In *ECML*, pages 318–329, 2006.
- [19] P. Nakov, L. Màrquez, A. Moschitti, W. Magdy, H. Mubarak, A. A. Freihat, J. Glass, and B. Randeree. SemEval-2016 task 3: Community question answering. In *Proceedings of SemEval '16*. ACL, 2016.
- [20] M. Nicosia, S. Filice, A. Barrón-Cedeño, I. Saleh, H. Mubarak, W. Gao, P. Nakov, G. Da San Martino, A. Moschitti, K. Darwish, L. Màrquez, S. Joty, and W. Magdy. QCRI: Answer selection for community question answering - experiments for Arabic and English. In *Proceedings of the 9th International Workshop on Semantic Evaluation*, SemEval '15, Denver, Colorado, USA, 2015.
- [21] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, ACL '02, pages 311–318, Philadelphia, Pennsylvania, USA, 2002.
- [22] R. Parker, D. Graff, K. Chen, J. Kong, and K. Maeda. *Arabic Gigaword Fifth Edition*. Linguistic Data Consortium (LDC), Philadelphia, 2011.
- [23] A. Pasha, M. Al-Badrashiny, M. Diab, A. E. Kholy, R. Eskander, N. Habash, M. Pooleery, O. Rambow, and R. Roth. Madamira: A fast, comprehensive tool for morphological analysis and disambiguation of arabic. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Reykjavik, Iceland, may 2014. European Language Resources Association (ELRA).
- [24] S. Petrov and D. Klein. Improved inference for unlexicalized parsing. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 404–411, Rochester, New York, April 2007. Association for Computational Linguistics.
- [25] A. Severyn and A. Moschitti. Structural relationships for large-scale learning of answer re-ranking. *SIGIR*, pages 741–750, 2012.
- [26] M. Snover, B. Dorr, R. Schwartz, L. Micciulla, and J. Makhoul. A study of translation edit rate with targeted human annotation. In *Proceedings of the 7th Biennial Conference of the Association for Machine Translation in the Americas*, AMTA '06, Cambridge, Massachusetts, USA, 2006.
- [27] K. Tymoshenko and A. Moschitti. Assessing the impact of syntactic and semantic structures for answer passages reranking. In *Proceedings of CIKM '15*, pages 1451–1460, New York, NY, USA, 2015. ACM.
- [28] K. Wang, Z. Ming, and T.-S. Chua. A syntactic tree matching approach to finding similar questions in community-based qa services. In *SIGIR*, pages 187–194, 2009.
- [29] M. Wise. Yap3: Improved detection of similarities in computer program and other texts. In *SIGCSE*, pages 130–134, 1996.
- [30] K. Zhang, W. Wu, H. Wu, Z. Li, and M. Zhou. Question retrieval with high quality answers in community question answering. In *CIKM*, pages 371–380, 2014.
- [31] G. Zhou, L. Cai, J. Zhao, and K. Liu. Phrase-based translation model for question retrieval in community question answer archives. In *ACL*, pages 653–662, 2011.