

# Cross-Language Question Re-Ranking

Giovanni Da San Martino, Salvatore Romeo, Alberto Barrón-Cedeño, Shafiq Joty, Lluís Màrquez,  
Alessandro Moschitti, Preslav Nakov  
Qatar Computing Research Institute, HBKU  
HBKU Research Complex, P.O. Box 5825  
Doha, Qatar  
{gmartino,sromeo,albarron,sjoty,lmarquez,amoschitti,pnakov}@hbku.edu.qa

## ABSTRACT

We study how to find relevant questions in community forums when the language of the new questions is different from that of the existing questions in the forum. In particular, we explore the Arabic–English language pair. We compare a kernel-based system with a feed-forward neural network in a scenario where a large parallel corpus is available for training a machine translation system, bilingual dictionaries, and cross-language word embeddings. We observe that both approaches degrade the performance of the system when working on the translated text, especially the kernel-based system, which depends heavily on a syntactic kernel. We address this issue using a cross-language tree kernel, which compares the original Arabic tree to the English trees of the related questions. We show that this kernel almost closes the performance gap with respect to the monolingual system. On the neural network side, we use the parallel corpus to train cross-language embeddings, which we then use to represent the Arabic input and the English related questions in the same space. The results also improve to close to those of the monolingual neural network. Overall, the kernel system shows a better performance compared to the neural network in all cases.

## CCS CONCEPTS

•Information systems → Learning to rank; Question answering; Multilingual and cross-lingual retrieval; Similarity measures;

## KEYWORDS

Community Question Answering; Cross-language Approaches; Question Retrieval; Kernel-based Methods; Neural Networks; Distributed Representations

## ACM Reference format:

Giovanni Da San Martino, Salvatore Romeo, Alberto Barrón-Cedeño, Shafiq Joty, Lluís Màrquez, Alessandro Moschitti, Preslav Nakov. 2017. Cross-Language Question Re-Ranking. In *Proceedings of SIGIR '17, Shinjuku, Tokyo, Japan, August 07–11, 2017*, 4 pages.  
DOI: <http://dx.doi.org/10.1145/3077136.3080743>

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 '17, Shinjuku, Tokyo, Japan*

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
978-1-4503-5022-8/17/08...\$15.00

DOI: <http://dx.doi.org/10.1145/3077136.3080743>

## 1 INTRODUCTION

In this paper, we study the problem of question re-ranking, which is an important task of the more general problem of community Question Answering (cQA). In particular, we address question re-ranking in a cross-language (CL) setting, i.e., where the language of the new question is different from the language of the candidate questions. We explore alternative ways to adapt kernel-based systems for English into this setting, when the query language is Arabic. This is an interesting scenario because state-of-the-art cQA models rely upon relational syntactic/semantic structures, using Tree Kernels (TKs) [9], and these might be difficult to port across translation-based models. We compare the kernel machines to feed-forward neural networks (FNN), which have been known to perform well for cQA [21].

We first explore a standard approach in CLIR: translating the input questions and applying our monolingual systems on the English-translated text. Our second approach, which is novel, is based on a CL TK—which does not require any translation—as it is applied directly to pairs of Arabic and English trees. This tree kernel makes use of a statistical bilingual dictionary extracted from a parallel corpus. The FNN system can also make use of the parallel corpus by learning cross-language embeddings, which we further use in order to compare the Arabic and the English input representations directly.

We tested our approaches on the benchmark datasets from the SemEval-2016 task 3 on cQA [22], which we enriched with Arabic new questions. The results show that machine translation does not drastically degrade the ranking performance, probably because of the robustness of our similarity features. Most importantly, the use of the cross-language tree kernels almost fills the gap with respect to the monolingual system.

## 2 RELATED WORK

Question re-ranking can be approached from several different angles. Cao et al. [6] tackled it by comparing representations based on topic term graphs, i.e., by judging topic similarity and question focus. Jeon, Croft, and Lee [15] and Zhou et al. [33] dodged the lexical gap between questions by assessing their similarity on the basis of a (monolingual) translation model. Wang et al. [31] computed a similarity function on the syntactic-tree representations of the questions. A different approach using topic modeling for question retrieval was introduced by Ji et al. [16] and Zhang et al. [32], who used LDA topic modeling to learn the latent semantic topics in order to retrieve similar questions. Dos Santos et al. [8] used neural networks for the same purpose.

Cross-language approaches have mainly focused on Question Answering (QA). This has been fostered by multiple challenges such as the Multilingual QA Challenge at CLEF 2008 [10], NTCIR-8's Advanced Cross-lingual Information Access (ACLIA) [20], and DARPA's Broad Operational Language Technologies (BOLT) IR task [26]. Usually, the full question is translated using an out-of-the-box system in order to address CL-QA [14, 17]. Ture and Boschee [29] proposed supervised models to combine different translation settings. Some approaches translate only keywords [24]. To the best of our knowledge, no research has been carried out on CL question re-ranking before. Regarding cross-language tree kernels, the only previous study relates to mapping natural language to an artificial language (SQL) [11, 12]. We use a similar cross-language tree kernel along with the new idea of deriving relational links [25] using cross-language dictionaries.

### 3 TASK AND CORPORA

We experiment with data from the SemEval-2016 Task 3 on Community Question Answering [22], which we further augment with translations as described below. We focus on subtask B, which targets *question-question similarity* (QS). Given a new question  $q_o$  and the set of ten related questions from the QatarLiving forum  $q_1, q_2, \dots, q_{10}$ , retrieved by a search engine, the goal is to re-rank the related questions according to their similarity with respect to the new question. The relationship between  $q_o$  and  $q_i, i \in \{1, 2, \dots, 10\}$ , is described with a label: PERFECTMATCH, RELEVANT, and IRRELEVANT. The goal is to rank the questions with the first two labels higher than those with the latter label. Note that the *questions* in this dataset are generally long multi-sentence stories which are written in informal English; full of typos and ungrammaticalities. The SemEval data has 267 new questions for training, 50 for development, and 70 for testing, and ten times as much  $q_o$ - $q_i$  pairs: 2,670, 500, and 700, respectively.

Based on this data, we simulated a *cross-language cQA setup*. We first got the 387 new train+dev+test questions translated into Arabic by professional translators.<sup>1</sup> Then, we used these Arabic versions of the questions as input with the goal of re-ranking the ten related English questions.

We also used an Arabic-English parallel corpus, which includes the publicly available TED and OPUS corpora [28]. We used this corpus in order to train an in-house phrase-based Arabic-English machine translation (MT) system,<sup>2</sup> and also to extract a bilingual dictionary in order to learn cross-language embeddings, as described in Sections 4 and 5 below.

### 4 A KERNEL-BASED SYSTEM

We address the re-ranking task described in Section 3 by using the scoring function of a binary ( $\{\text{PERFECTMATCH} \cup \text{RELEVANT}\}$  vs.  $\text{IRRELEVANT}$ ) classifier based on support vector machines (SVM):  $r(q_o, q_i) = \sum_j^n y_j \alpha_j K((q_o, q_i), (q_o^j, q_i^j))$  in order to rank all the related questions  $q_i$  with respect to their corresponding new question  $q_o$ . Here  $K(\cdot)$  is a kernel function assessing how similar two pairs of questions are. We use a combination of kernels on tree pairs and features as described below.

<sup>1</sup>The extension of the dataset is available at <http://alt.qcri.org/resources/cqa>.

<sup>2</sup>The MT system also uses a language model trained with the English Gigaword.

### 4.1 Tree Kernels

Given a pair of syntactic trees of the questions and a kernel for trees  $K^T$ , we define the following kernel applied to  $(q_o, q_i), (q_o^j, q_i^j)$ :

$$K^T(t(q_o, q_i), t(q_o^j, q_i^j)) + K^T(t(q_i, q_o), t(q_i^j, q_o^j)), \quad (1)$$

where  $t(x, x')$  is a string transformation method that returns the parse tree for the text  $x$ , further enriching it with *REL*ational tags computed with respect to the syntactic tree of  $x'$ . Typically, REL tags are assigned to the matching words between  $x$  and  $x'$ , and they are propagated to the parent and grand-parent nodes (i.e., up to 2 levels). This kernel in the monolingual setting is described in [4].

Note that Eq. (1) can be applied to pairs  $(q_o, q_i)$  in which  $q_o$  and  $q_i$  are texts in different languages, since in Eq. (1) the new (resp. related) questions are only compared to new (resp. related) questions: this produces the kernel space of tree fragment pairs as shown in [11, 12], where the pair members are in different languages. Moreover, the definition of  $t(x, x')$  is more complicated in case  $x$  and  $x'$  are in different languages as, in addition to using separate Arabic and English parsers, we need to define methods for matching words in different languages. Given the rich morphology of Arabic, this is not a trivial task.

**Cross-Language Tree Matching.** In order to match the lexical items from both trees, we created a word-level Arabic-to-English statistical dictionary using IBM's Model 1 [5] over our bilingual corpus, which we pre-processed using Farasa [7] in order to share the segmentation and diacritization of the Arabic syntactic parser.

### 4.2 Feature Vectors

We combined the above tree kernels linearly with RBF kernels applied to the following four feature vectors:

*ConvKNN features.* We used the 21 features proposed in [4] computing similarities between the new and related questions such as: longest common subsequences, Jaccard coefficient, word containment, cosine similarity. Since such similarities can be only computed when the two texts are in the same language, we use the English translation to obtain them for the cross-language system.

*Embedding features.* We used three types of vector-based embeddings in order to encode the text of a question: (1) **GOOGLE\_VECTORS**: 300-dimensional embedding vectors, pre-trained on Google News [19]; (2) **QL\_VECTORS**: we trained domain-specific vectors using **WORD2VEC** on all available QatarLiving data, both annotated and raw (as provided for SemEval-2016 Task 3). (3) **SYNTAX**: we parsed the questions using the Stanford neural parser [27], and we used the final 25-dimensional vector that is produced internally as a by-product of parsing. We did not use the embeddings themselves, but the cosines between the embeddings of a new and of a related question.

*MTE features.* We used the following MT evaluation metrics, which compare the similarity between the new and a related question as in [13]: (1) **BLEU**; (2) **NIST**; (3) **TER v0.7.25**; (4) **METEOR v1.4**; (5) **Unigram PRECISION**; (6) **Unigram RECALL**. We further used various components involved in the computation of BLEU, as features:  $n$ -gram precisions,  $n$ -gram matches, total number of  $n$ -grams ( $n=1,2,3,4$ ), lengths of the related and of the new questions, length ratio between them, and BLEU's brevity penalty.




Figure 1: Feed-forward neural network for QS.

*Task-specific features.* We computed various task-specific features, most of them introduced in the SemEval-2015 Task 3 on cQA [23]. This includes some question-level features: (1) number of URLs/images/emails/phone numbers; (2) number of tokens/sentences; (3) average number of tokens; (4) type/token ratio; (5) number of nouns/verbs/adjectives/adverbs/pronouns; (6) number of positive/negative smileys; (7) number of single/double/triple exclamation/interrogation symbols; (8) number of interrogative sentences (based on parsing); (9) number of words that are not in `WORD2VEC`'s Google News vocabulary. Also, some question-question pair features: (10) count ratio in terms of sentences/tokens/nouns/verbs/adjectives/adverbs/pronouns; (11) count ratio of words that are not in `WORD2VEC`'s Google News vocabulary. Finally, we also have one meta feature: (12) reciprocal rank of the related question.

## 5 A NEURAL NETWORK SYSTEM

Given the small size of the training set, we used a simple Feed-forward Neural Network (FNN), depicted in Figure 1. The input is a pair  $(q_o, q_i)$ . We map the input elements to fixed-length vectors  $(z_{q_o}, z_{q_i})$  using their *syntactic* and *semantic embeddings* (described in Section 4.2). The network then models the interactions between the input embeddings by passing them through two non-linear hidden layers (rectified linear units, ReLU). Additionally, the network also considers *pairwise features*  $\phi(q_o, q_i)$  between the two input elements that go directly to the output layer, and also through the second hidden layer. In our case,  $\phi(q_o, q_i)$  is the concatenation of the MTE and the task-specific features described in Section 4.2, which are also used by the kernel-based system. The following equations describe the transformation:  $\mathbf{h}_1 = f(U[z_{q_o}, z_{q_i}])$ ;  $\mathbf{h}_2 = f(V[\mathbf{h}_1, \phi(q_o, q_i)])$ , where  $U$  and  $V$  are the weight matrices in the first and in the second hidden layer.

The output layer of the neural network computes a sigmoid on the output layer weights and on the pairwise features in order to determine whether  $q_i$  is relevant with respect to the new question  $q_o$ . We train the models by minimizing the cross-entropy between the predicted distributions and the target distributions, i.e., the gold labels.

**Cross-language embeddings.** Using our parallel Arabic-English corpus, we trained cross-language embeddings using the bivec method by Luong et al. [18], a bilingual extension of word2vec, which has reported excellent results on semantic tasks close to ours [30]. Using these CL embeddings allows us to compare directly representations of the Arabic  $q_o$  and the English  $q_i$  input questions. In particular, we trained 200-dimensional word embeddings using the parameters described in [30], with a context window of size five and iterating for five epochs.

system	new question	MAP	
		dev.	test
1. IR rank	English	71.35	74.75
2. UH-PRHLT (SemEval)	English	—	76.70
3. ConvKN (SemEval)	English	—	76.02
4. SVM + TK	English	73.02	<b>77.41</b>
5. FNN	English	72.52	76.26
6. SVM + TK <sub>MT</sub>	Translated	72.94	76.67
7. SVM	Translated	71.99	76.36
8. FNN	Translated	72.44	75.73
9. SVM + TK <sub>CL</sub>	Arabic	73.34	<b>77.14</b>
10. FNN + CL_emb	Arabic	72.27	76.06

Table 1: MAP scores on the development and test datasets.

## 6 EXPERIMENTS

We consider three scenarios: (i) **Original**, i.e., the SemEval 2016 setup, (ii) **Translated**, in which  $q_o$  are originally in Arabic and machine-translated into English, and (iii) **Arabic**, in which  $q_o$  are in Arabic. In all settings, we apply the tree kernel in Eq. (1). However, we distinguish when the kernel is applied to the original English trees  $q_o$  (TK), the translated ones (TK<sub>MT</sub>), and the Arabic ones (TK<sub>CL</sub>). In all experiments, we kept the default parameters for the kernels and we selected the  $C$  parameter of SVM on the development set, trying {0.01, 0.1, 1, 10}. For the FNN models, we used the development set for early stopping based on MAP as well as for parameter optimization.

Table 1 shows Mean Average Precision (MAP) on the development and on the test datasets. The first block contains the reference results on the original English test set (rows 1–5). *IR rank* corresponds to the Google-generated ranking, which is a hard-to-beat baseline. *UH-PRHLT* and *ConvKN* are the two best-performing systems at SemEval-2016 Task 3 (see [22] for details). *SVM + TK* is the kernel-based system presented in this paper, which reproduces *ConvKN* and adds our extra features (cf. Section 4.2). Finally, *FNN* is the neural network model presented in this paper.

Our SVM model (row 4) shows a sizable improvement over *ConvKN* on the test set, which means that our extra features are strong. Actually, the SVM results are also better than the best system at SemEval-2016 (+0.71 MAP). The FNN model shows also competitive performance, but below the SVM system (-1.15 MAP). The monolingual result from SVM (77.41 MAP) is the upper bound performance when considering the results in the CL scenario.

The second block (rows 6–8) shows the results of our systems in the “translated” setting. One concern about applying the TK approach in this setting was that the translated text might be grammatically broken and the parser could produce low-quality parse trees. Still, the SVM system degrades its performance only to 76.67 MAP when applied after machine-translating the Arabic new questions (i.e., -0.74 MAP points below the upper bound).

Row 7 shows the result for SVM without TK. Its MAP score is 76.36, which is slightly below the previous score of 76.67 obtained with TK; this shows that the two kernels provide complementary information. Comparatively, the FNN model degrades the performance even less when working with the translated Arabic query (75.73, row 8 vs. 76.26, row 5). This indicates again that the features we use are robust to translation.

The last block in the table (rows 9-10) shows the results of the systems using the CL kernels and representations. The SVM system scores 77.14 MAP when using the CL kernel. This is above the results with  $TK_{MT}$ . This final MAP value is very close to the upper bound system (77.41). In conclusion, achieving a similar ranking quality to that of the monolingual setting is possible by departing from Arabic text and using the novel cross-language tree kernel together with a robust feature set computed on the translated texts. Finally, the FNN system achieves slightly improved results when we add the input representation based on cross-language embeddings (row 10), reaching a MAP score, 76.06, that is very close to the monolingual FNN (76.26).

## 7 CONCLUSIONS AND FUTURE WORK

We studied the task of cross-language question re-ranking in community question answering. We first explored the possibility of using MT for translating an Arabic query question and then applying an English monolingual system. The results of two alternative systems for question re-ranking (kernel- and FNN-based) show a relatively small degradation in performance with respect to the monolingual setting on a well-established SemEval dataset. Furthermore, we showed that the performance gap in the SVM system can be almost closed by using a novel cross-language tree kernel, which compares directly the source and the target language trees. A cross-language input representation can also help the FNN system to close the gap with respect to the monolingual case. Finally, the performance of the SVM system is always superior to that of the FNN system in our setting. We conjecture that this is due to the relatively small size of the training set, and due to the information provided by the tree kernel (relations between syntactic sub-structures).

Our work enables interesting future research lines, e.g., (i) designing more accurate cross-language TKs using better Arabic structures and cross-language word matching and embeddings, (ii) combining the SVM and the FNN models, and (iii) exploring how far a system can go without machine translation.

## ACKNOWLEDGMENTS

This research was performed by the Arabic Language Technologies group at Qatar Computing Research Institute, HBKU, within the Interactive sYstems for Answer Search project (IYas).

## REFERENCES

- [1] 2008. *Evaluating Systems for Multilingual and Multimodal Information Access, 9th Workshop of CLEF 2008, Revised Selected Papers*. Aarhus, Denmark.
- [2] 2010. *Proc. of NTCIR-8 Workshop Meeting*. Tokyo, Japan.
- [3] 2016. *Proc. of the 10th SemEval Workshop*. San Diego, California, USA.
- [4] Alberto Barrón-Cedeño, Giovanni Da San Martino, Shafiq Joty, Alessandro Moschitti, Fahad Al-Obaidli, Salvatore Romeo, Kateryna Tymoshenko, and Antonio Uva. 2016. ConvKN at SemEval-2016 Task 3: Answer and Question Selection for Question Answering on Arabic and English Fora, See [3], 896–903.
- [5] Peter F. Brown, John Cocke, Stephen A. Della Pietra, Vicent J. Della Pietra, Frederick Jelinek, John D. Lafferty, Robert L. Mercer, and Paul S. Roossin. 1990. A Statistical Approach to Machine Translation. *Computational Linguistics* 16, 2 (1990), 79–85.
- [6] Yunbo Cao, Huizhong Duan, Chin-Yew Lin, Yong Yu, and Hsiao-Wuen Hon. 2008. Recommending Questions Using the Mdl-based Tree Cut Model. In *Proceedings of WWW '08*. New York, New York, USA, 81–90.
- [7] Kareem Darwish and Hamdy Mubarak. 2016. Farasa: A New Fast and Accurate Arabic Word Segmenter. In *Proceedings of LREC*. Portorož, Slovenia.
- [8] Cicero Dos Santos, Luciano Barbosa, Dasha Bogdanova, and Bianca Zadrozny. 2015. Learning Hybrid Representations to Retrieve Semantically Equivalent Questions. In *Proceedings of ACL '15*. Beijing, China, 694–699.
- [9] Simone Filice, Danilo Croce, Alessandro Moschitti, and Roberto Basili. 2016. KeLP at SemEval-2016 Task 3: Learning Semantic Relations between Questions and Answers, See [3], 1116–1123.
- [10] Pamela Forner, Anselmo Peñas, Eneko Agirre, Iñaki Alegria, Corina Forascu, Nicolas Moreau, Petya Osenova, Prokopis Prokopidis, Paulo Rocha, Bogdan Sacaleanu, Richard F. E. Sutcliffe, and Erik F. Tjong Kim Sang. 2008. Overview of the CLEF 2008 Multilingual Question Answering Track, See [1], 262–295.
- [11] Alessandra Giordani and Alessandro Moschitti. 2010. Semantic Mapping Between Natural Language Questions and SQL Queries via Syntactic Pairing. In *Proceedings of NLDB '09*. Saarbrücken, Germany, 207–221.
- [12] Alessandra Giordani and Alessandro Moschitti. 2012. Translating Questions to SQL Queries with Generative Parsers Discriminatively Reranked. In *Proceedings of COLING '12*. Mumbai, India, 401–410.
- [13] Francisco Guzmán, Lluís Márquez, and Preslav Nakov. 2016. Machine Translation Evaluation Meets Community Question Answering. In *Proceedings of ACL '16*. Berlin, Germany, 460–466.
- [14] Sven Hartrumpf, Ingo Glöckner, and Johannes Leveling. 2008. University of Hagen at QA@CLEF 2008: Efficient Question Answering with Question Decomposition and Multiple Answer Streams, See [1], 421–428.
- [15] Jiwoon Jeon, W. Bruce Croft, and Joon Ho Lee. Finding Similar Questions in Large Question and Answer Archives. In *Proceedings of CIKM '05*. Bremen, Germany, 84–90.
- [16] Zongcheng Ji, Fei Xu, Bin Wang, and Ben He. 2012. Question-Answer Topic Model for Question Retrieval in Community Question Answering. In *Proceedings of CIKM '12*. Maui, Hawaii, USA, 2471–2474.
- [17] Chuan-Jie Lin and Yu-Min Kuo. 2010. Description of the NTOU Complex QA System, See [2], 47–54.
- [18] Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Bilingual Word Representations with Monolingual Quality in Mind. In *Proceedings of the 1st Workshop on Vector Space Modeling for NLP*. Denver, Colorado, USA, 151–159.
- [19] Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013. Linguistic Regularities in Continuous Space Word Representations. In *Proceedings of NAACL-HLT '13*. Atlanta, Georgia, USA, 746–751.
- [20] Teruko Mitamura, Hideki Shima, Tetsuya Sakai, Noriko Kando, Tatsunori Mori, Koichi Takeda, Ruihua Lin, Chin-Yew Song, Chuan-Jie Lin, and Cheng-Wei Lee. 2010. Overview of the NTCIR-8 ACLIA Tasks: Advanced Cross-Lingual Information Access, See [2], 15–24.
- [21] Preslav Nakov, Lluís Márquez, and Francisco Guzmán. 2016. It Takes Three to Tango: Triangulation Approach to Answer Ranking in Community Question Answering. In *Proceedings of EMNLP '16*. Austin, Texas, USA, 1586–1597.
- [22] Preslav Nakov, Lluís Márquez, Alessandro Moschitti, Walid Magdy, Hamdy Mubarak, abed Alhakim Freihat, Jim Glass, and Bilal Randeree. 2016. SemEval-2016 Task 3: Community Question Answering, See [3], 525–545.
- [23] Massimo Nicosia, Simone Filice, Alberto Barrón-Cedeño, Iman Saleh, Hamdy Mubarak, Wei Gao, Preslav Nakov, Giovanni Da San Martino, Alessandro Moschitti, Kareem Darwish, Lluís Márquez, Shafiq Joty, and Walid Magdy. QCR: Answer Selection for Community Question Answering - Experiments for Arabic and English. In *Proceedings of SemEval '15*. Denver, Colorado, USA, 203–209.
- [24] Han Ren, Donghong Ji, and Jing Wan. 2010. WHU Question Answering System at NTCIR-8 ACLIA Task, See [2], 31–36.
- [25] Aliaksei Severyn and Alessandro Moschitti. 2012. Structural Relationships for Large-scale Learning of Answer Re-Ranking. In *Proceedings of SIGIR '12*. Portland, Oregon, USA, 741–750.
- [26] Ian Soboroff, Kira Griffith, and Stephanie Strassel. 2016. The BOLT IR Test Collections of Multilingual Passage Retrieval from Discussion Forums. In *Proceedings of SIGIR '16*. New York, New York, USA, 713–716.
- [27] Richard Socher, John Bauer, Christopher D. Manning, and Ng Andrew Y. 2013. Parsing with Compositional Vector Grammars. In *Proceedings of ACL '13*. Sofia, Bulgaria, 455–465.
- [28] Jörg Tiedemann. 2012. Parallel Data, Tools and Interfaces in OPUS. In *Proceedings of LREC '12*. Istanbul, Turkey.
- [29] Ferhan Ture and Elizabeth Boschee. 2016. Learning to Translate for Multilingual Question Answering. In *Proceedings of EMNLP '16*. Austin, Texas, USA, 573–584.
- [30] Shyam Upadhyay, Manaal Faruqui, Chris Dyer, and Dan Roth. 2016. Cross-lingual Models of Word Embeddings: An Empirical Comparison. In *Proceedings of ACL*. Berlin, Germany, 1661–1670.
- [31] Kai Wang, Zhaoyan Ming, and Tat-Seng Chua. 2009. A Syntactic Tree Matching Approach to Finding Similar Questions in Community-based QA Services. In *Proceedings of SIGIR '09*. Boston, Massachusetts, USA, 187–194.
- [32] Kai Zhang, Wei Wu, Haocheng Wu, Zhoujun Li, and Ming Zhou. 2014. Question Retrieval with High Quality Answers in Community Question Answering. In *Proceedings of CIKM '14*. Shanghai, China, 371–380.
- [33] Guangyou Zhou, Li Cai, Jun Zhao, and Kang Liu. 2011. Phrase-based translation model for question retrieval in community question answer archives. In *Proceedings of ACL*. Portland, Oregon, USA, 653–662.