

Multitask Learning with Deep Neural Networks for Community Question Answering

Daniele Bonadiman*
Università di Trento

Antonio Uva**
Università di Trento

Alessandro Moschitti†
Amazon

In this paper, we developed a deep neural network (DNN) that learns to solve simultaneously the three tasks of the cQA challenge proposed by the SemEval-2016 Task 3, i.e., question-comment similarity, question-question similarity and new question-comment similarity. The latter is the main task, which can exploit the previous two for achieving better results. Our DNN is trained jointly on all the three cQA tasks and learns to encode questions and comments into a single vector representation shared across the multiple tasks. The results on the official challenge test set show that our approach produces higher accuracy and faster convergence rates than the individual neural networks. Additionally, our method, which does not use any manual feature engineering, approaches the state of the art established with methods that make heavy use of it.

1. Introduction

Community Question Answering (cQA) websites enable users to freely ask questions in web forums and expect some good answers in the form of comments from the other users. Given the large number of question/answer pairs available on cQA sites, researchers started to investigate the possibility to exploit user-generated content for training automatic QA systems. Unfortunately, the text involved in the cQA scenario is rather noisy, therefore, providing models that outperform the simple bag-of-words representation can result rather difficult. The challenge, SemEval-2016 Task 3 “Community Question Answering”, has been designed to study the above problems: the participants were supposed to build a fully automatic system for cQA. In particular, given a fresh user question, q_{new} , and a set of forum questions, Q , answered by a comment set, C , the main task consists of determining whether a comment $c \in C$ is a pertinent answer of q_{new} or not. This task can be divided into three sub-tasks:

- (A) predict if a comment produced in response to a question contains a valid answer;

* Dept. of Information Engineering and Computer Science (DISI) - Via Sommarive, 9, 38123 Povo, Trento, Italy. E-mail: d.bonadiman@unitn.it

** Dept. of Information Engineering and Computer Science (DISI) - Via Sommarive, 9, 38123 Povo, Trento, Italy. E-mail: antonio.uva@unitn.it

† Manhattan Beach, CA, USA, 90266. E-mail: amosch@amazon.com. This work was carried out when the author was at the University of Trento.

- (B) re-rank a set of questions according to their relevancy with respect to the original question; and
- (C) predict if a comment produced in response to a previous question posed on the cQA forum represents a valid answer to a fresh question.

Traditionally, these tasks have been tackled by designing systems/classifiers that target each task separately. Each classifier accepts a vector encoding a text pair (e.g., a question/question or a question/answer pair) in input by using many complex lexical syntactic or semantic features and, then, computing similarity between these representations. However, this approach suffers from the drawbacks of requiring a “customized” set of features for each task being solved.

Recent work on deep neural networks (DNNs) for Multitask Learning (MTL) (Collobert and Weston 2008; Liu et al. 2015) showed that is possible to *jointly train* a general system that solves different tasks simultaneously. Inspired by the success of MTL, in this paper, we propose a DNN model that leverages the data from the three cQA tasks of SemEval. Indeed, as the three tasks are highly related, we claim that cQA can benefit from this approach. We show that, despite the fact that does not require any feature engineering, our DNN can approach the performance of the best systems, which use heavy feature engineering. Additionally, we are going to make the corpora for studying MTL on this interesting challenge available to the research community.

2. cQA Tasks at SemEval

The research problem issued by SemEval-2016 Task 3 is exemplified by Fig. 2: given a new question q_{new} , Task C is about directly retrieving a relevant comment from the entire community. This can also be achieved by solving Task B, which finds a similar question, q_{rel} , and then executing Task A, which selects good comments, c_{rel} , for q_{rel} . It should be noted that Task A classifies comments, specifically written by the users for q_{rel} , whereas Task C classifies comments written by the users for other, sometimes, similar questions. This means, it needs to filter out comments that can be partially related to q_{new} (because they correctly answer the related question, q_{rel}) but still not correctly answering q_{new} . Clearly, Task C classifier needs to tackle a much more semantically challenging task. Thus, tasks A and C are semantically and computationally rather different and together with Task B: they constitute an interesting MTL problem since differences and correlations are played at a very high semantic level.

2.1 Task A: Question-Comment Similarity

Given a question, q_{rel} , and its first 10 comments, c_{rel} , in the question threads, rerank the comments according to their relevance to q_{rel} . Relevancy is defined according to three classes: (i) *good*: the comment is definitively relevant; (ii) *potentially useful*: the comment is not good, but it still contains related information worth checking; and (iii) *bad*: the comment is irrelevant (e.g., it is part of a dialogue or unrelated to the topic). For evaluation purposes, both *potentially useful* and *bad* comments were considered as *bad*.

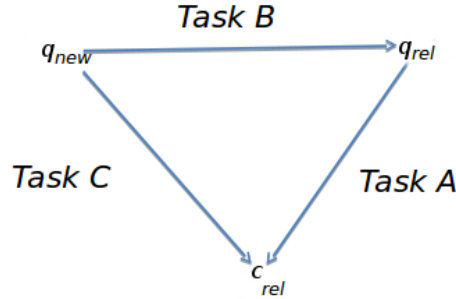


Figure 1

The 3 tasks of cQA at SemEval: the arrows show the relations between the original and the related questions and the related comments.

2.2 Task B: Question-Question Similarity

Given a new question, q_{new} , and its first 10 related questions (retrieved by a search engine), q_{rel} , rerank them according to their similarity with respect to q_{new} . Relevancy is expressed by three classes: (i) *perfect match*: the new and forum questions request roughly the same information, (ii) *relevant*: the new and forum questions ask for similar information, or (iii) *irrelevant*: the new and forum questions are completely unrelated. For evaluation purposes, both *perfect match* and *relevant* forum questions are considered as *relevant*.

2.3 Task C: New Question-Comment Similarity

Given a new question, q_{new} , and its first 10 related questions (retrieved by a search engine), q_{rel} , each associated with its first 10 comments, c_{rel} , appearing in its thread, rerank the 100 comments (10 questions \times 10 comments) according to their relevance with respect to q_{new} . Relevancy is defined similarly to task A.

2.4 Dataset

The data for the above-mentioned tasks is distributed in three datasets: train, dev and test sets. The distribution of questions and comments in each dataset varies across the different tasks: Task A contains 6,938 related questions and 40,288 comments. Each comment in the dataset was annotated with a label indicating its relevancy with respect to the related question. Task B contains 317 original questions. For each original question, 10 related questions were retrieved, summing to 3,169 related questions. Also in this case, the related questions were annotated with a relevancy label, which tells if they are relevant with respect to the user original question. Task C contains 317 original questions, together with 3,169 related questions (same as in Task B) and 31,690 comments. Each comment was labelled with its relevancy with respect to the original question.

3. A General Deep Architecture for cQA

All the previous tasks are about reranking questions or comments with respect to an original question. In the following, we describe a general architecture for solving them.

3.1 Deep Architecture for relational learning from pairs of text

A traditional approach to cQA is to learn a different classifiers for solving each of these three tasks, independently. For example, first a classifier can be trained to rerank a set of related questions retrieved by a search engine, using their similarity with respect to the user question (Task B). Then, another classifier can be trained to rerank the list of comments appearing in the threads of similar questions (Task A). Each of these classifiers uses a different set of task-dependent features. In this work, we use a neural network architecture for classifying text pairs. The network is fed using the different pairs, (q_{rel}, c_{rel}) , (q_{new}, q_{rel}) and (q_{new}, c_{rel}) , to learn the tasks A, B and C, respectively, and produces a similarity score that can be used to rerank questions or comments.

It is composed of two main components: (i) two sentence encoders that map input sentences i into fixed size vectors $x_{s_i} \in \mathbb{R}^m$, and (ii) a feed forward neural network that computes the similarity between these two sentence vectors.

The sentence encoders are composed of (i) a sentence matrix $s_i \in \mathbb{R}^{d \times |i|}$, where d is the size of the word embeddings, obtained by concatenating the vectors of the corresponding words in the input sentence $w_j \in s_i$, and (ii) a sentence model $f : \mathbb{R}^{d \times |i|} \rightarrow \mathbb{R}^m$, which maps the sentence matrix to a fixed size sentence embedding $x_{s_i} \in \mathbb{R}^m$.

The choice of the sentence model plays a crucial role as the resulting intermediate representation of the input sentences affects the successive steps of computing their similarity. Previous work in this direction uses different types of sentence models such as LSTM, distributional sentence model (average of word vectors), and convolutional sentence model. In particular, the latter is composed of a sequence of convolution and pooling feature maps, which have achieved the state of the art in various NLP tasks (Kalchbrenner, Grefenstette, and Blunsom 2014; Kim 2014).

In this paper, we used a CNN sentence model generated by a convolutional operation followed by a k -max pooling layer with $k = 1$, since it provides comparable performance to the LSTM on the task of new question-comment similarity, as shown in Table 2. The sentence encoder, $x_{s_i} = f(s_i)$, outputs a fixed-size embedding of the input sentence s_i . The sentence vectors, x_{s_i} , are concatenated together and given in input to a Multi-Layer Perceptron, which is constituted by a non-linear hidden layer and an sigmoid output layer.

3.2 Injecting Relational Information

All the tasks we consider require to model relations between words present in the two pieces of text. For this purpose, we encode the relation in forms of discrete features, as described in (Collobert et al. 2011), i.e., using an additional embedding layer. They augmented the word embedding with the corresponding feature embedding. Thus, given a word, w_j , its final word embedding is defined as $w_j \in \mathbb{R}^d$, where $d = d_w + d_{feat}$, where d_w is the size of the word embedding and d_{feat} is the size of the feature embedding.

We use a discrete feature, represented with an embedding of 5 dimensions, to encode matches between two words in the two input pieces of text. In particular, we associate each word w in the input sentences with a *word overlap* index $o \in \{0, 1\}$, where $o = 1$ means that w is shared by both Q and C (or by the two questions for task

B), i.e., overlaps, $o = 0$ otherwise. It should be noted that the embeddings described here cannot be considered as task specific features, manually handcrafted. They are part of the network, serve the purpose of injecting relational information between the representations of the two input texts and can be generally applied to different domains, data and tasks.

3.3 Adding the rank features

The SemEval problems concern reranking text initially ranked by Google and made available to the participants for tasks B and C. Considering that the Google rank is computed using powerful algorithms and a lot of resources, it is essential to encode it in our networks. There are several methods to achieve this. After some experiments, we opted for discretizing the rank values in 5 different bins of different sizes, i.e. $[1 - 2]$, $[2 - 5]$, $[5 - 10]$, $[10 - 25]$, $[25 - \infty]$. The rank feature is added to the joint layer, where the output of the sentence model is concatenated, using a table lookup operation. It should be noted that for each task, we use a different relation feature (overlap embeddings) between each pair of text.

4. MTL for cQA

MTL aims at learning several related tasks at the same time to improve some (or possibly all) tasks using joint information (Caruana 1997). MTL is particularly well suited for modeling Task C as it is a composition of tasks A and B, thus, it can benefit from having both questions q_{new} and q_{rel} as input to better model the interaction between the new question and the comment. More precisely, it can use the triplet $\langle q_{new}, q_{rel}, c_{rel} \rangle$ in the learning process, where the interaction between the triplet members is exploited during the joint training of the three models of the tasks, A, B and C. In fact, an improvement on question-comment similarity or on question-question similarity can lead to an improvement in the task of new question-comment similarity (Task C).

Additionally, each thread in the SemEval dataset is annotated with the labels for all the three tasks and therefore it is possible to apply joint learning directly.

4.1 Joint Learning Architecture

Our Joint learning architecture is depicted in Figure 2, it is a direct extension of the architecture proposed for Task C (Section 3.3). It takes the three sentences as input, i.e. a new question, q_{new} , the related question, q_{rel} , and its comment, c_{rel} , and produces three fixed size representations, $x_{q_{new}}$, $x_{q_{rel}}$ and $x_{c_{rel}}$, respectively.

These three representations are then concatenated ($h_j = [x_{q_{new}}, x_{q_{rel}}, x_{c_{rel}}]$) and fed to a hidden layer to create a shared representation of the input for the three tasks, $h_s = Wh_j$.

The output of this layer, h_s is then fed to three independent Multilayer Perceptrons (MLP) that produce the scores for the three tasks. To directly apply MTL, we use the binary cross-entropy instead of the max margin loss as our objective function. The main motivation is that such function is computed based on pairs of positive-negative examples that cannot be created with multiple labels. At training time, for each example, the loss is calculated on the three outputs of the network. The final loss is then the sum of the individual losses for the three tasks.

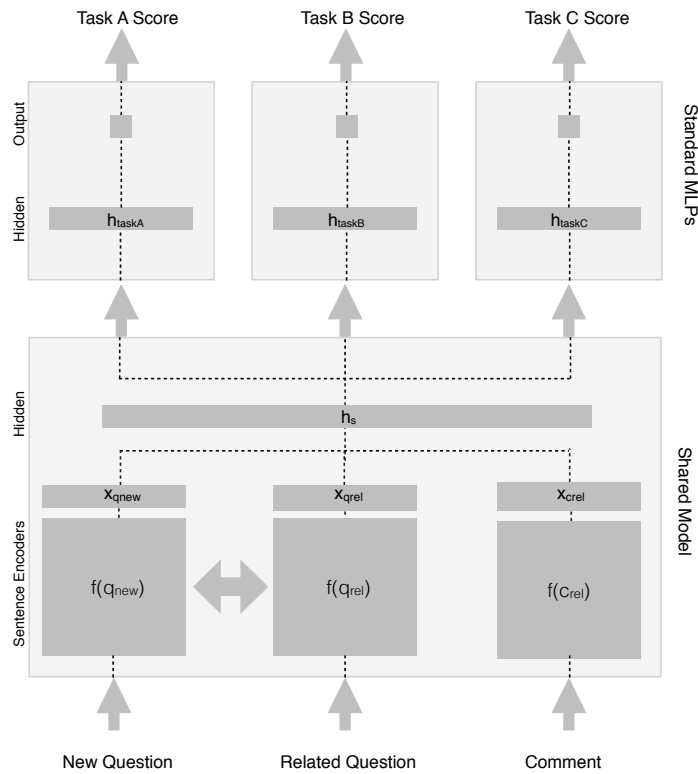


Figure 2 Our MTL architecture, where the three sentences are at the bottom. They pass through the sentence encoders. The output is concatenated and fed to a hidden layer whose output is passed to three independent multi-layer perceptrons, which produce the scores for the individual tasks. The double arrow, \leftrightarrow , indicates a shared sentence model between q_{new} and q_{rel} .

4.2 Shared Sentence Models

The SemEval dataset contains ten times less new questions than related questions by construction. However, all questions, q_{new} included, are supposed to be of the same nature. Thus we can certainly use a shared text model for modeling better representations for both new and related questions. Formally, let $x_d = f(d, \theta)$ be a sentence model for document d with parameters θ , i.e., the embedding weights and the convolutional filters. In our original formulation, each sentence model uses a different set of parameters $\theta_{q_{new}}$, $\theta_{q_{rel}}$ and $\theta_{c_{rel}}$. However, for the question representation, we also used the same set of parameters θ_q . Such shared sentence model is illustrated by a double arrow in Figure 2.

5. Experiments

5.1 Setup

We encode input sentences with fixed-sized vectors using a convolutional operation of size 5 and a k -max pooling operation with $k = 1$, i.e., similarly to (Severyn and Moschitti

Table 1

Percentage of positive examples in the training datasets for each task.

	Task A	Task B	Task C
Train	37.51%	39.41%	9.9%
Train + ED	37.47%	64.38%	21.25%

2015, 2016). We use two non-linear hidden layers (with hyperbolic tangent activation, Tanh), whose size is equal to the size of the previous layer, i.e., the join layer. We include information such as word overlaps and rank position as embedding with an additional lookup table with vectors of size $d_{feat} = 5$.

Pre-processing: both questions and comments are tokenized and lowercased (to reduce the dimensionality of the dictionary and therefore of the embedding matrix). Moreover, question subject and body are concatenated to create a unique question. For computational reasons, we opted to limit the size of the input text at 100 words: we did not observe any degradation in performance.

Word Embeddings: for all the proposed models, we pre-initialize the word embedding matrices with standard skipgram embedding of dimensionality 50 trained on the English Wikipedia dump using word2vec toolkit (Mikolov et al. 2013).

Training: The network is trained using SGD with shuffled mini-batches using the rmsprop update rule (Tieleman and Hinton 2012). The model learns until the validation loss stops improving, with patience $p = 10$, i.e., the number of epochs to wait before early stopping, if no progress on the validation set is obtained. In fact, early stopping (Prechelt 1998) allows us to avoid overfitting and improving the generalization capabilities of the network. For the MTL architecture, we employed two different stopping criteria. The first is to stop training when the global validation loss does not improve anymore (the sum of the individual losses of the three tasks). The second, instead, saves three different models and evaluates them when the individual losses stop improving. Since the three tasks converge at different epochs, the first method may lead to sub-optimal results for the individual tasks, but only one model is needed at test time.

To improve generalization and avoid co-adaptation of features, we opted for adding dropout (Srivastava et al. 2014) between all the layers of the network. We experimented with different dropout rates (0.2, 0.4) for the inputs and (0.3, 0.5, 0.7) for the hidden layers obtaining better results with the highest values, i.e., 0.4 and 0.7.

Dataset: Table 1 reports the labels distributions on the train dataset. It is important to note that the dataset for Task C presents a higher number of negative than positive examples. For this reason, we automatically extended the training dataset (ED) with new positive matches for tasks B and C, respectively. This process is done by creating the (q_{rel}, c_{rel}) pairs for each q_{rel} from the training set for Task A and creating triples of the form $(q_{rel}, q_{rel}, c_{rel})$, where the label for question-question similarity is obviously positive and the labels for Task C are inherited from those of Task A. The resulting dataset contains 34,100 triples and its relevance label distribution is presented in the last row of Table 1. The extended version of the dataset with the annotation for MTL is made available for download for comparison purposes ¹.

¹ <https://ikernels-portal.disi.unitn.it/repository/>

Table 2
Impact of CNN vs. LSTM sentence models on the baseline network for Task C.

Model	MAP	MRR
LSTM	43.91	49.28
CNN	44.43	49.01
CNN Train	44.43	49.01
CNN Train + ED ²	44.77	52.07

Table 3
Results on the validation and test set for the proposed models

Models	Task A: question-comment similarity				Task B: question-question similarity				Task C: new question-comment similarity			
	DEV		TEST		DEV		TEST		DEV		TEST	
	MAP	MRR	MAP	MRR	MAP	MRR	MAP	MRR	MAP	MRR	MAP	MRR
Random	-	-	59.53	67.83	-	-	46.98	50.96	-	-	15.01	15.19
IR Baseline	-	-	52.80	58.71	-	-	74.75	83.79	-	-	40.36	45.83
Kelp	-	-	79.19	86.42	-	-	-	-	-	-	-	-
UH-PRHLT	-	-	-	-	-	-	76.70	83.02	-	-	-	-
SUper-team	-	-	-	-	-	-	-	-	-	-	55.41	61.48
$\langle q_{rel}, c_{rel} \rangle$	68.93	76.46	74.73	81.18	-	-	-	-	-	-	-	-
$\langle q_{new}, q_{rel} \rangle$	-	-	-	-	74.19	83.26	73.70	82.13	-	-	-	-
$\langle q_{new}, c_{rel} \rangle$	-	-	-	-	-	-	-	-	44.77	52.07	41.95	47.21
$\langle q_{new}, q_{rel}, c_{rel} \rangle$	-	-	-	-	-	-	-	-	45.59	51.04	46.99	55.64
$\langle q_{new}, q_{rel}, c_{rel} \rangle + \leftrightarrow$	70.69	77.19	75.52	82.11	72.92	80.20	72.88	80.58	47.82	53.03	46.45	51.72
MTL (BC)	-	-	-	-	74.22	80.40	73.68	81.59	47.80	52.31	48.58	55.77
MTL (AC)	70.11	76.50	75.43	82.46	-	-	-	-	46.34	51.54	48.49	54.01
MTL (ABC)	69.93	76.27	74.42	81.68	70.68	75.85	71.07	80.11	49.63	55.47	49.87	55.73
MTL (ABC)*	70.70	77.48	74.89	81.80	74.21	81.93	72.23	80.33	49.63	55.47	49.87	55.73
MTL (weighted score)	-	-	-	-	-	-	-	-	-	-	52.67	55.68

Measures: we report the results of our models in terms of MAP and MRR. Both provide a higher score if the relevant items are higher in the rank. However, MAP takes into account the rank of all of the relevant items with respect to the irrelevant ones. MRR only considers the first relevant retrieved item with respect to all the others.

5.2 Impact of the sentence models

Table 2 shows a comparison between CNN and LSTM sentence models when used in our general architecture (see Sec. 3) for solving Task C. We derived the results from the development set³. We observe that the two sentence models show comparable results. For the rest of the experiments, we used the CNN sentence model, since it shows faster convergence rate and more stable results with respect to the LSTM sentence model. In the second part of Table 2, we demonstrate that using the extended dataset for solving Task C leads to higher results than the original one. In particular, we noted that there is an improvement of 3 points in MRR.

² Extended Dataset for Task C computed using questions from Task A.

³ In this work, the dataset Train-part2 were used as development set.

5.3 Results of individual models

Table 3 shows the results of our individual and MTL models, in comparison with the Random and Information Retrieval baselines of the challenge (first grouped row), and the three-top systems of SemEval 2016, KeLP (Filice et al. 2016), UH-PRHLT (Franco-Salvador et al. 2016), SUper-team (Mihaylova et al. 2016) (second grouped row).

The third grouped row shows the performance of the individual models when trained on input pairs, $\langle q_{rel}, c_{rel} \rangle$, $\langle q_{new}, q_{rel} \rangle$ and $\langle q_{new}, c_{rel} \rangle$ for task A, B and C, respectively. The model for the three tasks is the same (described in Sec. 3). These results show that the individual models can generalize well enough on all tasks. In particular, on Task B, they achieve the best results of our proposed model (the numbers in bold indicate the best results among the proposed models).

The fourth grouped row illustrates the models exploiting the joint input, $\langle q_{new}, q_{rel}, c_{rel} \rangle$, but no joint learning is carried out, i.e., the networks for the different tasks are trained individually. The results show that a small degradation of performance happens in Task B, while Task A slightly improves. These variations may be due to the fact that tasks A and B can be efficiently solved using the standard pairwise approach, thus the extra text introduced in the model may just add some noise. However, using the shared sentence model for q_{new} and q_{rel} of the tasks B and C (indicated with \leftrightarrow) improves the overall performance.

5.4 Results of MTL models

The shared input representation shows good results on all tasks, thus, in the last set of experiments, we jointly trained (i) tasks B and C, (ii) tasks A and C and finally (iii) the three tasks together.

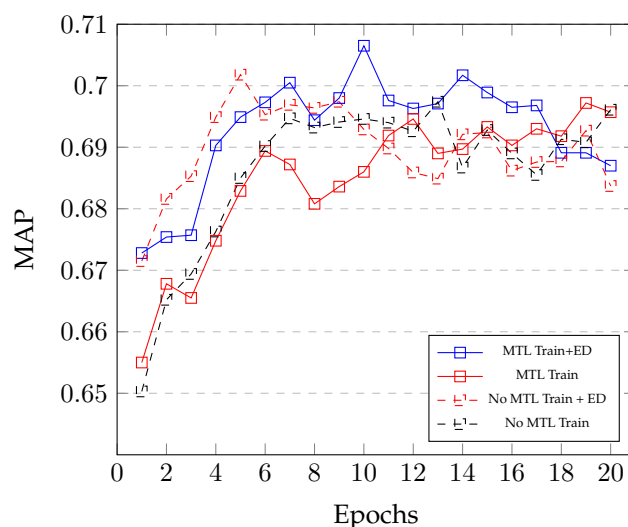
The results are reported in the fifth grouped row. It is interesting to note that the major boost in terms of performance is obtained when we jointly train all the three tasks. In fact, the MTL architecture improves the individual model in terms of MAP by about 2 absolute points on the DEV set and by 3 absolute points on the TEST set for Task C, while the performance on the other tasks tends to degrade. However, if the three different models are evaluated at different epochs of training, e.g., see MTL(ABC)*, it is possible to obtain accuracy comparable to the individual models for all the three tasks. As previously explained, when applying MTL, the individual objective functions converge at different epochs. Therefore, when the global loss reaches the minimum, it is possible that individual models are sub-optimal.

Indeed, the comparison between the learning curves (on the development set) for Task B (Figure 4) and Task C (Figure 5) shows that for the former, models achieve earlier convergence rate (epoch 2) while for the latter they converge later (epoch 16). Moreover, Figure 3 shows that the results on Task A are not badly affected by jointly training models with the other two tasks.

Finally, the learning curves show that our networks trained in MTL tend to have faster convergence rate than the individual models: this is a very interesting result.

We also experimented with shallower networks and SVMs using the prediction scores from the different classifiers in a stacking approach, and obtained results far below the baselines⁴.

⁴ We did not include these results as they do not provide interesting findings.

**Figure 3**

Learning curves for Task A on the dev. set; dotted and solid lines represent the individual and multi-task models, respectively.

Comparison with the state of the art. Our models would have ranked 4th on Task C of the Semeval 2016 competition ⁵, i.e., the main task of the challenge. In contrast, our models for the other two tasks, which do not benefit from the overall MTL architecture would have achieved a middle position (8th). These results are important since our proposed MTL architecture obtains a placement very close to the top system, without requiring task-specific features, which in cQA are extremely important, e.g., the thread-level features.

Finally, one reason of why we do not achieve the state of the art on Task C is due to the difference between training and test data. Several challenge participants solved this problem by using a weighted sum between the score of the Task A classifier and the Google rank as a strong features for modeling Task C. We followed a similar approach estimating the weight MTL on the dev set and using the computed score to rank the comments of the test set. This improved the MAP of our MTL by about 2.8 absolute points on the test set, obtaining a result comparable with the model ranked 2nd on Task C at the Semeval 2016 competition.

6. Related Work

Previous work related to the topics presented in this paper spans three major research areas: Question Retrieval (targeting question similarity), Passage Reranking (targeting question and answer similarity) and MTL. In the following, we will report the most important works in these areas.

⁵ <http://alt.qcri.org/semeval2016/task3/index.php?id=results>

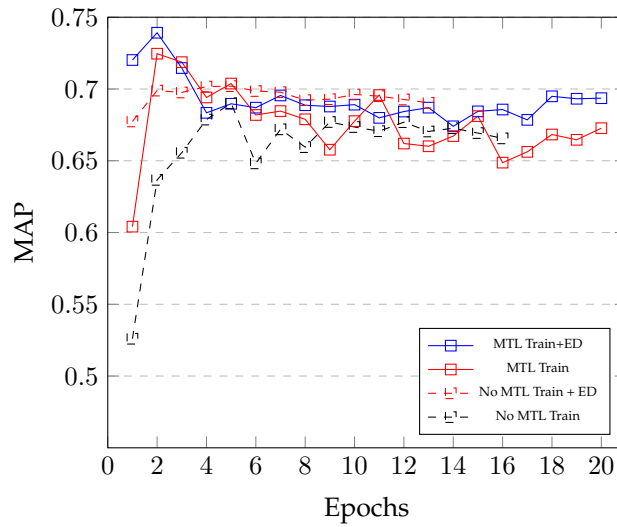


Figure 4 Learning curves for Task B on the development set; dotted lines represent the individual models, while the solid lines represent the multi-task ones.

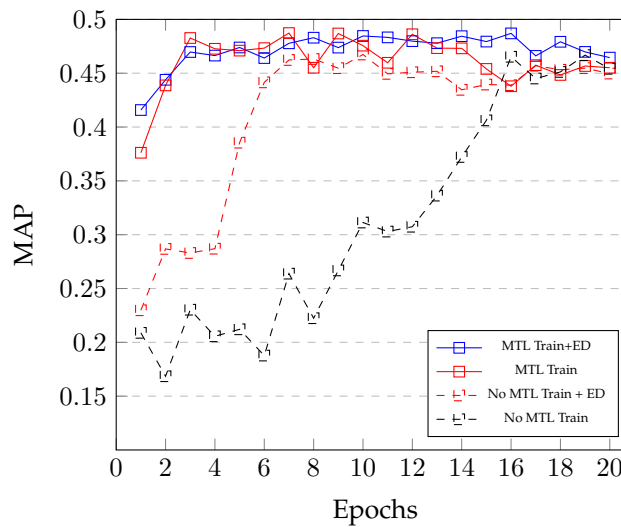


Figure 5 Learning curves for Task C on the development set; dotted lines represent the individual models, while the solid lines represent the multi-task ones.

Question-Question Similarity. Determining question similarity remains one of the main tasks needed to be solved in cQA due to difficult problems such as “lexical gap”. Early approaches on question similarity used statistical machine translation techniques to measure similarity between questions. For example, [Jeon, Croft, and Lee 2005] and [Zhou et al. 2011] used a language models based on word or phrase translation probabilities to estimate similarity between questions. However, effective approaches based

on statistical machine translation require lots of data to estimate word probabilities. Language models for question-question similarity were also explored by [Cao et al. 2009]. These models exploit information from the category structure of Yahoo! Answers when computing similarity between two questions. Instead, [Duan et al. 2008] propose an approach that identifies the topic and focus in a text and compute similarity between two input questions by matching the extracted topic and focus information. A different approach to question-question similarity is provided by [Ji et al. 2012] and [Zhang et al. 2014]. They use LDA to learn the probability distribution over the topics that generate the question/answers pairs. Later, this distribution is used to measure similarity between questions.

Question-Answer Similarity. In recent years, many models have been proposed for computing similarity of an answer with respect to a question. For example, [Yao et al. 2013] trained a conditional random field based on a set of powerful features, such as tree-edit distance between question and answer trees: these also enable the extraction of answers from pre-retrieved sentences. [Heilman and Smith 2010] use a linear classifier using syntactic features to solve different tasks such as recognizing textual entailment, paraphrases and answer selection. [Wang, Smith, and Mitamura 2007] propose the use of Quasi-synchronous grammars to select short answers for TREC questions. This is done by learning syntactic and semantic transformation from the question to the answer trees. [Wang and Manning 2010] propose a probabilistic Tree-Edit model with structured latent variables for solving textual entailment and question answering. An advanced model based on structural representations was proposed in (Moschitti et al. 2007; Moschitti 2008; Severyn and Moschitti 2012; Severyn, Nicosia, and Moschitti 2013; Severyn and Moschitti 2013, 2015; Tymoshenko and Moschitti 2015). These model use SVM with kernels to learn structural patterns between questions and answers encoded in form of shallow syntactic parse trees.

Finally, [Wang and Nyberg 2015] trained a long short-term memory model for selecting answers to TREC questions. Their model takes words from question and answer sentences as input and returns a score measuring the relevancy of an answer with respect to a given question. A recent work close to ours is (Guzmán, Márquez, and Nakov 2016), where the authors build a neural network for solving Task A of SemEval. However, this does not approach the problem as MTL.

Related work on MTL. A good overview on MTL, i.e., learning to solve multiple tasks by using a shared representation with mutual benefit, is given in (Caruana 1997). [Collobert and Weston 2008] trained a convolutional NN with MTL which, given an input sentence, performs many sequence labeling tasks. They showed that jointly training their system on different tasks, such as speech tagging, named entity recognition, etc., significantly improves the performance on the main task, i.e., semantic role labeling, without requiring hand-engineered features.

[Liu et al. 2015] is the most close work to ours. They used multi-task deep neural networks to map queries and documents into semantic vector representations. This representation is later used into two tasks: query classification and question-answer reranking. The results showed a competitive gain over strong baselines. In our work, we have presented an architecture that can also exploit joint representation of question and comments, given the strong interdependencies among the different SemEval Tasks.

7. Conclusion

In this paper we proposed several Deep Neural Networks for the task of automatic cQA. Our main result is a network that can effectively exploit the characteristics of the cQA task to carry out interesting MTL. Our network designed and trained in an MTL setting shows better accuracy and a higher convergence rate than the models independently trained. The results show that our MTL model approaches the performance of the models participating at the SemEval 2016 cQA competition. It should be noted that all the other challenge systems use domain specific features, which are both very important but also rather costly to engineer.

In the future, we would like to use more effective features and combine them with other machine learning methods.

References

- Cao, Xin, Gao Cong, Bin Cui, Christian Søndergaard Jensen, and Ce Zhang. 2009. The use of categorization information in language models for question retrieval. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM '09*, pages 265–274, New York, NY, USA. ACM.
- Caruana, Rich. 1997. Multitask learning. *Machine Learning*, 28(1):41–75, Jul.
- Collobert, Ronan and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pages 160–167, New York, NY, USA. ACM.
- Collobert, Ronan, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research*, 12:2493–2537.
- Duan, Huizhong, Yunbo Cao, Chin-Yew Lin, and Yong Yu. 2008. Searching questions by identifying question topic and question focus. In *Proceedings of ACL-08: HLT*, pages 156–164. Association for Computational Linguistics.
- Filice, Simone, Danilo Croce, Alessandro Moschitti, and Roberto Basili. 2016. Kelp at semeval-2016 task 3: Learning semantic relations between questions and answers. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1116–1123. Association for Computational Linguistics.
- Franco-Salvador, Marc, Sudipta Kar, Tamar Solorio, and Paolo Rosso. 2016. Uh-prhlt at semeval-2016 task 3: Combining lexical and semantic-based features for community question answering. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 814–821. Association for Computational Linguistics.
- Guzmán, Francisco, Lluís Màrquez, and Preslav Nakov. 2016. Machine translation evaluation meets community question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 460–466. Association for Computational Linguistics.
- Heilman, Michael and Noah A. Smith. 2010. Tree edit models for recognizing textual entailments, paraphrases, and answers to questions. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1011–1019. Association for Computational Linguistics.
- Jeon, Jiwoon, W. Bruce Croft, and Joon Ho Lee. 2005. Finding similar questions in large question and answer archives. In *Proceedings of the 14th ACM International Conference on Information and Knowledge Management, CIKM*.
- Ji, Zongcheng, Fei Xu, Bin Wang, and Ben He. 2012. Question-answer topic model for question retrieval in community question answering. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management, CIKM '12*, pages 2471–2474, New York, NY, USA. ACM.
- Kalchbrenner, Nal, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 655–665. Association for Computational Linguistics.

- Kim, Yoon. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751. Association for Computational Linguistics.
- Liu, Xiaodong, Jianfeng Gao, Xiaodong He, Li Deng, Kevin Duh, and Ye-Yi Wang. 2015. Representation learning using multi-task deep neural networks for semantic classification and information retrieval. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 912–921. Association for Computational Linguistics.
- Mihaylova, Tsvetomila, Pepa Gencheva, Martin Boyanov, Ivana Yovcheva, Todor Mihaylov, Momchil Hardalov, Yasen Kiprova, Daniel Balchev, Ivan Koychev, Preslav Nakov, Ivelina Nikolova, and Galia Angelova. 2016. Super team at semeval-2016 task 3: Building a feature-rich system for community question answering. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 836–843. Association for Computational Linguistics.
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In Chris J.C. Burges, Léon Bottou, Max Welling, Zoubin Ghahramani, and Iliyan Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- Moschitti, Alessandro. 2008. Kernel methods, syntax and semantics for relational text categorization. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management, CIKM 2008*, pages 253–262, Napa Valley, California, USA, October 26–30.
- Moschitti, Alessandro, Silvia Quarteroni, Roberto Basili, and Suresh Manandhar. 2007. Exploiting syntactic and shallow semantic kernels for question answer classification. In *ACL 2007, Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics*, June 23–30, 2007, Prague, Czech Republic.
- Prechelt, Lutz. 1998. Early stopping-but when? In *Neural Networks: Tricks of the trade*. Springer, pages 55–69.
- Severyn, Aliaksei and Alessandro Moschitti. 2012. Structural relationships for large-scale learning of answer re-ranking. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 741–750. ACM.
- Severyn, Aliaksei and Alessandro Moschitti. 2013. Automatic feature engineering for answer selection and extraction. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP '13*, pages 458–467, Seattle, Washington, USA.
- Severyn, Aliaksei and Alessandro Moschitti. 2015. Learning to rank short text pairs with convolutional deep neural networks. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 373–382. ACM.
- Severyn, Aliaksei and Alessandro Moschitti. 2016. Modeling relational information in question-answer pairs with convolutional neural networks. *CoRR*, abs/1604.01178.
- Severyn, Aliaksei, Massimo Nicosia, and Alessandro Moschitti. 2013. Building structures from classifiers for passage reranking. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, pages 969–978. ACM.
- Srivastava, Nitish, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Tieleman, Tijmen and Geoffrey Hinton. 2012. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4.
- Tymoshenko, Kateryna and Alessandro Moschitti. 2015. Assessing the impact of syntactic and semantic structures for answer passages reranking. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management, CIKM 2015*, pages 1451–1460, Melbourne, VIC, Australia, October 19 - 23.
- Wang, Di and Eric Nyberg. 2015. A long short-term memory model for answer sentence selection in 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 707–712. Association for Computational Linguistics.
- Wang, Mengqiu and Christopher Manning. 2010. Probabilistic tree-edit models with structured latent variables for textual entailment and question answering. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 1164–1172. Coling 2010 Organizing Committee.

- Wang, Mengqiu, Noah A. Smith, and Teruko Mitamura. 2007. What is the jeopardy model? a quasi-synchronous grammar for qa. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*.
- Yao, Xuchen, Benjamin Van Durme, Chris Callison-Burch, and Peter Clark. 2013. Answer extraction as sequence tagging with tree edit distance. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 858–867. Association for Computational Linguistics.
- Zhang, Kai, Wei Wu, Haocheng Wu, Zhoujun Li, and Ming Zhou. 2014. Question retrieval with high quality answers in community question answering. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM '14*, pages 371–380, New York, NY, USA. ACM.
- Zhou, Guangyou, Li Cai, Jun Zhao, and Kang Liu. 2011. Phrase-based translation model for question retrieval in community question answer archives. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 653–662. Association for Computational Linguistics.