APPLYING DEEP LEARNING TO ANSWER SELECTION: A STUDY AND AN OPEN TASK

Minwei Feng, Bing Xiang, Michael R. Glass, Lidan Wang, Bowen Zhou

IBM Watson Yorktown Heights, NY, USA, 10598 <mfeng|bingxia|mrglass|wangli|zhou>@us.ibm.com

ABSTRACT

We apply a general deep learning framework to address the non-factoid question answering task. Our approach does not rely on any linguistic tools and can be applied to different languages or domains. Various architectures are presented and compared. We create and release a QA corpus and setup a new QA task in the insurance domain. Experimental results demonstrate superior performance compared to the baseline methods and various technologies give further improvements. For this highly challenging task, the top-1 accuracy can reach up to 65.3% on a test set, which indicates a great potential for practical use.

Index Terms— Answer Selection, Question Answering, Convolutional Neural Network (CNN), Deep Learning, Spoken Question Answering System

1. INTRODUCTION

Natural language understanding based spoken dialog system has been a popular topic in the past years of artificial intelligence renaissance. Many of those influential systems include a question answering module, e.g. Apple's Siri, IBM's Watson and Amazon's Echo. In this paper, we address the Question Answering (QA) module in those spoken QA systems. We treat the QA from a text matching and selection perspective. IBM's Watson system [1] is a classical example of the traditional way of doing Question Answering (QA). In this work we utilize a deep learning framework to accomplish the answer selection which is a key step in the QA task. Hence QA is studied from an answer matching and selection perspective. Given a question q and an answer candidate pool $\{a_1, a_2, ..., a_s\}$ for that question (s is the pool size), the goal is to find the best answer candidate $a_k, 1 \leq k \leq s$. If the selected answer a_k is inside the ground truth set (one question could have more than one correct answer) of question q, the question q is considered to be answered correctly, otherwise it is answered incorrectly. From the definition, the QA problem can be regarded as a binary classification problem. For each question, for each answer candidate, it may be appropriate or not. In order to find the best pair, we need a metric to measure

the matching degree of each QA pair so that the QA pair with highest metric value will be chosen.

The above definition is general. The only assumption made is that for every question there is an answer candidate pool. In practice, the pool can be easily built by using a general search engine like Google Search or an information retrieval software library like Apache Lucene.

We created a data set by collecting question and answer pairs from the internet. All these question and answer pairs are in the insurance domain. The construction of this insurance domain QA corpus was driven by the intense scientific and commercial interest in this domain. We released this corpus¹ to create an open QA task, enabling other researchers to utilize it and supporting a fair comparison among different methods. The corpus consists of four parts: train, development, test1 and test2. Table 1 gives the data statistics. All experiments conducted in this paper are based on this corpus. To our best knowledge, it is the first time an insurance domain QA task has been released.

Our QA task requires specifying an answer candidate pool for each question in the development, test1 and test2 parts of the corpus. The released corpus contains totally 24 981 unique answers. It is possible to use the whole answer space as the candidate pool, so that each question must be compared with 24 981 answer candidates. However, this is impractical due to time consuming computations. In this paper, we set the pool size to be 500, so that it is both practical and still a challenging task. We put the ground truth answers into the pool and randomly sample negative answers from the answer space until the pool size reaches 500.

The technology described in this paper with the released data set and benchmark task is targeting potential applications like online customer service. Hence it is not supposed to handle question answering tasks that require reasoning, e.g. is tomorrow Tuesday? (answer depends on if today is Monday.) The rest of the paper is organized as follows: Sec. 2 describes the different architectures used this work; Sec. 3 provides the experimental setup details; experimental results and discussions are presented in Sec. 4; Sec. 5 contains related work

¹git clone https://github.com/shuzi/insuranceQA.git

	Questions	Answers	Question Word Count
Train	12887	18540	92095
Dev	1000	1454	7158
Test1	1800	2616	12893
Test2	1800	2593	12905

Table 1: Corpus statistics: first two columns are the question and answer quantity; notice there could be multiple answers for some questions so that the answer quantity is larger than the question quantity; third column is the question total word count. The total number of answers is 24 981 and the whole answer text contains 2 386 749 words

and finally we draw conclusions in Sec. 6.

2. MODEL DESCRIPTION

In this section we describe the proposed deep learning framework and many variations based on that framework. However, the main idea of those different systems is the same: learn a distributed vector representation of a given question and its answer candidates and then use a similarity metric to measure the matching degree. We first developed two baseline systems for comparison.

2.1. Baseline Systems

The first baseline system is a bag-of-words model. Step one is to train a word embedding by [2]. This word embedding provides the word vector for each token in the question and its candidate answer. From these, the baseline system produces the idf-weighted sum of word vectors for the question and for all of its answer candidates. This produces a vector representation for the question and each answer candidate. The last step is to calculate the cosine similarity between each question/candidate pair. The pair with highest cosine similarity is returned as the answer. The second baseline is an information retrieval (IR) baseline. The state-of-the-art weighted dependency model (WD) [3, 4] is used. The WD model employs a weighted combination of term-based and term proximitybased ranking features to score each candidate answer. Example features include counts of question bigrams in ordered and unordered windows of different sizes in each candidate answer, in addition to simple unigram counts. The basic idea is that important bigrams or unigrams in the question should receive higher weights when their frequencies are computed. Thus, the feature weights are assigned in accordance to the importance of the question unigrams or bigrams that they are defined over, where the importance factor is learned as part of the model training process. Row 1 and 2 (first column Idx) of Table 2 are the baseline system results.

2.2. CNNs-based System

In this paper, a QA framework based on Convolutional Neural Networks (CNN) is presented. As summarized in Chapter 11 of [5], a CNN leverages three important ideas that can help improve a machine learning system: sparse interaction, parameter sharing and equivariant representation. Sparse interaction contrasts with traditional neural networks where each output is interactive with each input. In a CNN, the filter size (or kernel size) is usually much smaller than the input size. As a result, the output is only interactive with a narrow window of the input. Parameter sharing refers to reusing the filter parameters in the convolution operations, while the element in the weight matrix of traditional neural network will be used only once to calculate the output. Equivariant representation is related to the idea of k-MaxPooling which is usually combined with a CNN. In this paper we always set k = 1. So each filter of the CNN represents some feature, and after the convolution operation, the 1-MaxPooling value represents the highest degree that the input contains the feature. The position of that feature in the input is irrelevant due to the convolution. This property is very useful for many NLP applications. Below is an example to demonstrate our CNN implementation.

$$\begin{pmatrix} w_{11} & w_{21} & w_{31} & w_{41} \\ w_{12} & w_{22} & w_{32} & w_{42} \\ w_{13} & w_{23} & w_{33} & w_{43} \end{pmatrix} \bigodot \begin{pmatrix} f_{11} & f_{21} \\ f_{12} & f_{22} \\ f_{13} & f_{23} \end{pmatrix}$$
(1)

The left matrix W is the input sentence. Each word is represented by a 3-dimensional word embedding vector and the input length is 4. The right matrix F represents the filter. The 2-gram filter size is 3×2 . The convolution output of the input W and the filter F is a 3-dim vector O, assuming zero padding has been done so that only a narrow convolution is conducted.

$$o_{1} = w_{11}f_{11} + w_{12}f_{12} + w_{13}f_{13} + w_{21}f_{21} + w_{22}f_{22} + w_{23}f_{23}$$

$$o_{2} = w_{21}f_{11} + w_{22}f_{12} + w_{23}f_{13} + w_{31}f_{21} + w_{32}f_{22} + w_{33}f_{23}$$

$$o_{3} = w_{31}f_{11} + w_{32}f_{12} + w_{33}f_{13} + w_{41}f_{21} + w_{42}f_{22} + w_{43}f_{23}$$
(2)

After 1-MaxPooling, the maximum of the 3 values will be kept for the filter F which indicates the highest degree that filter F matches the input W.

2.3. Training and Loss Function

Different architectures will be described later. However all those different architectures share the same training and testing mechanism. In this paper we minimize a ranking loss similar to [6] [7]. During training, for each training question Q there is a positive answer A^+ (the ground truth). A training instance is constructed by pairing this A^+ with a negative answer A^- (a wrong answer) sampled from the whole answer space. The deep learning framework generates vector representations for the question and the two candidates:



Fig. 1: Architecture I. Q for question; A for answer; P is 1-MaxPooling; T is tanh layer; HL for hidden layer and HL already includes tanh as its activation function.



Fig. 2: Architecture II . QA means the weights of corresponding layer are shared by Q and A .

 V_Q , V_{A^+} and V_{A^-} . The cosine similarities $\cos(V_Q,V_{A^+})$ and $\cos(V_Q,V_{A^-})$ are calculated and the distance between the two similarities is compared to a margin: $\cos(V_Q,V_{A^+}) - \cos(V_Q,V_{A^-}) < m$. m is the margin. When this condition is satisfied, the implication is that the vector space embedding either ranks the positive answer below the negative answer, or does not sufficiently rank the positive answer above the negative answer. If $\cos(V_Q,V_{A^+}) - \cos(V_Q,V_{A^-}) >= m$ there is no update to the parameters and a new negative example is sampled until the margin is less than m (to reduce running time we set maximum 50 times in this paper). The hinge loss function is hence defined as follows:

$$L = max \{0, m - cos(V_Q, V_{A^+}) + cos(V_Q, V_{A^-})\}$$
(3)

For testing, we calculate the $cos(V_Q, V_{candidate})$ between the question Q and each answer candidate $V_{candidate}$ in the pool (size 500). The candidate answer with largest cosine similarity is selected.

2.4. Architectures

In this subsection we demonstrate several proposed architectures for this QA task. Figure 1 shows the Architecture I. Q is the input question provided as input to the first hidden layer HL_Q. The hidden layer (HL) is defined as z = tanh(Wx+B). W is the weight matrix; B is the bias vector; x is input; z is the output of the activation function tanh. The output then flows to the CNN layer CNN_Q, applied to extract question side features. P is the MaxPooling layer (we always use 1-MaxPooling in this paper) and T is the tanh layer. Similar to the question side, the answer A is processed by HL_A and then features are extracted by CNN_A. 1-MaxPooling P



Fig. 3: Architecture III . HL for hidden layer. Add another HL_Q and HL_A after CNN_{QA} .



Fig. 4: Architecture IV . Add another shared hidden layer HL_{QA} after CNN_{QA} .

and *tanh* layer T will function in the last step. The result is a vector representation for both question and answer. The final output is the cosine similarity between these vectors. Row 3 of Table 2 is the Architecture I result.

Figure 2 is the Architecture II. The main difference compared to Architecture I is that both question and answer sides share the same HL and CNN weights. Row 4 of Table 2 is the Architecture II result.

We also consider architectures with a hidden layer after the CNN. Figure 3 is the Architecture III in which another HL_Q is added at the question side after the CNN and another HL_A is added at the answer side after the CNN. Row 5 of Table 2 is the Architecture III result. Architecture IV, shown in Figure 4, is similar except the second HL of both question and answer share the same HL_{QA} weights. The rows 6 and 7 of Table 2 are the Architecture IV results.

Figure 5 is the Architecture V where two layers of CNN_{QA} are deployed. In section 2.2 we show the convolution output is a vector (3-dim in that example). This is true only for CNNs with a single filter. By applying multiple filters the result is a matrix. If there are 4 filters utilized for the example in section 2.2, the output is the following matrix.

$$\begin{pmatrix}
o_{11} & o_{21} & o_{31} \\
o_{12} & o_{22} & o_{32} \\
o_{13} & o_{23} & o_{33} \\
o_{14} & o_{24} & o_{34}
\end{pmatrix}$$
(4)

Each row represents the output of one filter and each column represents a bigram of the input. This matrix is the input to the next CNN_{QA} layer. For this second layer, every bigram is effectively one "word" and the previous filter's output for that bigram is its word embedding. Row 11 of Table 2 is the result for Architecture V. Architecture VI in Figure 6 is similar

Idx	Dev	Test1	Test2	Description
1	31.9	32.1	32.2	Baseline: Bag-of-words
2	52.7	55.1	50.8	Baseline: metzler-bendersky IR model
3	44.2	41.7	39.5	Architecture I: $HL_Q(200) HL_A(200) CNN_Q(1000) CNN_A(1000)$ 1-MaxPooling Tanh
4	58.2	57.8	53.6	Architecture II: HL _{QA} (200) CNN _{QA} (1000)1-MaxPooling Tanh
5	36.1	33.6	32.7	Architecture III: HL _{QA} (200) CNN _{QA} (1000) HL _Q (1000) HL _A (1000) 1-MaxPooling Tanh
6	51.4	50.5	46.1	Architecture IV: HL _{QA} (200) CNN _{QA} (1000) HL _{QA} (1000) 1-MaxPooling Tanh
7	47.0	46.7	43.0	Architecture IV: HL _{QA} (200) CNN _{QA} (1000) HL _{QA} (500) 1-MaxPooling Tanh
8	60.6	59.2	55.1	Architecture II: HL _{QA} (200) CNN _{QA} (2000) 1-MaxPooling Tanh
9	61.5	61.3	57.8	Architecture II: HL _{QA} (200) CNN _{QA} (3000) 1-MaxPooling Tanh
10	61.8	62.8	59.2	Architecture II: $HL_{QA}(200) CNN_{QA}(4000)$ 1-MaxPooling Tanh (best result in this table)
11	59.7	59.3	55.6	Architecture V: HL _{QA} (200) CNN _{QA} (1000) CNN _{QA} (1000) 1-MaxPooling Tanh
12	59.9	60.6	55.9	Architecture VI: HL _{QA} (200) CNN _{QA} (1000) CNN _{QA} (1000) 1-MaxPooling Tanh (2COST)
13	59.9	58.7	53.8	Architecture II: HL _{QA} (200) Augmented-CNN _{QA} (1000) 1-MaxPooling Tanh
14	60.0	60.3	54.3	Architecture II: HL _{QA} (200) Augmented-CNN _{QA} (2000) 1-MaxPooling Tanh
15	61.7	62.2	56.3	Architecture II: HL _{QA} (200) Augmented-CNN _{QA} (3000) 1-MaxPooling Tanh

Table 2: Experimental Results. HL(200) means the hidden layer size is 200; CNN(1000) means there are 1000 filters used; top one precision of Dev, Test1 and Test2 have been reported.

to Architecture V except we utilize layer-wise supervision. After each CNN_{QA} layer there is 1-MaxPooling and a tanh layer so that the cost function can be calculated and back-propagation can be conducted. The result of Architecture VI is in row 12 of Table 2.

We have tried another three techniques to improve Architecture II in Figure 2. First, the CNN filter quantity has been increased, see row 8 9 and 10 of Table 2. Second, the convolution operation has been augmented to include skip-bigrams. Consider the example in section 2.2, for the input and one filter in Eq. 1, the augmented convolution operation will not only produce Eq. 2 but also the following discontinuous convolution:

$$o_4 = w_{11}f_{11} + w_{12}f_{12} + w_{13}f_{13} + w_{31}f_{21} + w_{32}f_{22} + w_{33}f_{23}
o_5 = w_{21}f_{11} + w_{22}f_{12} + w_{23}f_{13} + w_{41}f_{21} + w_{42}f_{22} + w_{43}f_{23}$$
(5)

The 1-MaxPooling will still be applied to get the largest value among $[o_1, o_2, o_3, o_4, o_5]$ so that this filter is automatically adapted to match a bigram or skip-bigram feature. Rows 13 14 and 15 of Table 2 show the results. Third, we investigate the similarity metric. Until now, we have been using the cosine similarity which is widely adopted for vector space models. However, is cosine similarity the best option for this task? Table 3 is the results for similarity metric study. Some metrics include hyperparameters and experiments with various hyperparameters have been conducted. We propose two novel metrics (GESD and AESD) which demonstrate superior performance.

3. EXPERIMENTAL SETUP

The deep learning framework in this paper has been built from scratch using Java. To improve speed, we adopt the HOG-WILD approach [8]. Each thread processes one training in-



Fig. 5: Architecture V. Two shared CNN_{QA} .



Fig. 6: Architecture VI. Two shared CNNQA. Two cost functions.

stance at one time and updates the weights of the neural networks. There is no locking in any thread. The word embedding (100 dimensions) is trained by word2vec [2] and used for initialization. Word embeddings are also parameters and are optimized for the QA task. Stochastic Gradient Descent is the optimization strategy and the L2-norm is also added in the loss function. In this paper, the weight of the L2-norm is 0.0001, the learning rate is 0.01 and margin m is 0.009. Those hyperparameters are chosen based on previous experiences in using deep learning on this data and they are not very sensitive within reasonable range. The utilized computing re-

Dev	Test1	Test2	Description
58.2	57.8	53.6	cosine: $k(x,y) = \frac{xy^{T}}{\ x\ \ y\ }$
58.5	57.1	53.3	polynomial: $k(x, y) = (\gamma x y^{T} + c)^d, \ \gamma = 0.5, \ d = 2, \ c = 1$
56.8	54.6	52.6	polynomial: $k(x, y) = (\gamma x y^{T} + c)^d, \ \gamma = 1.0, \ d = 2, \ c = 1$
55.0	53.6	48.2	polynomial: $k(x, y) = (\gamma x y^{T} + c)^d, \ \gamma = 1.5, d = 2, c = 1$
57.1	53.7	51.5	polynomial: $k(x, y) = (\gamma x y^{T} + c)^d, \ \gamma = 0.5, d = 3, c = 1$
55.3	52.4	48.7	polynomial: $k(x, y) = (\gamma x y^{T} + c)^d, \ \gamma = 1.0, \ d = 3, \ c = 1$
52.5	51.0	47.2	polynomial: $k(x, y) = (\gamma x y^{T} + c)^d, \ \gamma = 1.5, \ d = 3, \ c = 1$
61.3	59.9	57.0	sigmoid: $k(x, y) = tanh(\gamma xy^{T} + c), \gamma = 0.5, c = 1$
61.6	60.2	57.1	sigmoid: $k(x, y) = tanh(\gamma xy^{T} + c), \gamma = 1.0, c = 1$
60.2	60.2	55.7	sigmoid: $k(x, y) = tanh(\gamma xy^{\intercal} + c), \gamma = 1.5, c = 1$
60.0	60.3	54.7	RBF: $k(x, y) = exp(-\gamma x - y ^2), \gamma = 0.5$
60.2	57.0	54.4	RBF: $k(x, y) = exp(-\gamma x - y _{x}^{2}), \gamma = 1.0$
58.4	57.3	53.8	RBF: $k(x, y) = exp(-\gamma x - y ^2), \gamma = 1.5$
60.8	60.3	57.0	euclidean: $k(x, y) = \frac{1}{1 + x - y }$
42.2	42.5	38.2	exponential: $k(x, y) = exp(-\gamma x - y _1), \gamma = 0.5$
41.4	39.5	36.0	exponential: $k(x, y) = exp(-\gamma x - y _1), \gamma = 1.0$
48.2	45.1	41.6	exponential: $k(x, y) = exp(-\gamma x - y _1), \gamma = 1.5$
51.0	49.5	46.4	manhattan: $k(x, y) = \frac{1}{1 + \ x - y\ _1}$
62.5	61.4	59.0	GESD: $k(x, y) = \frac{1}{1+\ x-y\ } \cdot \frac{1}{1+exp(-\gamma(xy^{\intercal}+c))}, \gamma = 0.5, c = 1$
62.9	62.1	59.3	GESD: $k(x, y) = \frac{1}{1+ x-y } \cdot \frac{1}{1+exp(-\gamma(xy^{\intercal}+c))}, \gamma = 1.0, c = 1$
62.6	62.1	59.2	GESD: $k(x,y) = \frac{1}{1+\ x-y\ } \cdot \frac{1}{1+exp(-\gamma(xy^{T}+c))}, \gamma = 1.5, c = 1$
63.1	61.9	58.2	AESD: $k(x, y) = \frac{10.5}{1+ x-y } + \frac{10.5}{1+exp(-\gamma(xy^{T}+c))}, \gamma = 0.5, c = 1$
63.4	61.7	58.7	AESD: $k(x, y) = \frac{0.5}{1 + x - y } + \frac{1 + e^{-\gamma (x - \gamma (x + c))}}{1 + e^{-\gamma (x - \gamma (x + c))}}, \gamma = 1.0, c = 1$
62.8	62.0	57.7	AESD: $k(x, y) = \frac{1 + \left\ x - y\right\ }{1 + \left\ x - y\right\ } + \frac{1 + \left\ x - y\right\ }{0.5} + \frac{1 + \left\ x - y\right\ }{0.5}, \gamma = 1.5, c = 1$
63.5	62.5	60.2	GESD: $k(x, y) = \frac{1}{1+ x-y } \cdot \frac{1}{1+exp(-\gamma(xyT+c))}, \gamma = 1.0, 2000$ filters
64.3	65.1	61.0	GESD: $k(x, y) = \frac{1}{1+ x-y } \cdot \frac{1}{1+exp(-\gamma(xy+c))}, \gamma = 1.0, 3000 \text{ filters}$
65.4	65.3	61.0	GESD: $k(x, y) = \frac{1}{1+ x-y } \cdot \frac{1}{1+exp(-\gamma(xyT+c))}, \gamma = 1.0, 4000 \text{ filters}$
64.5	62.7	60.1	AESD: $k(x, y) = \frac{0.5}{1+ x-y } + \frac{1+\exp(-\gamma(xy+e))}{1+\exp(-\gamma(xy+e))}, \gamma = 1.0, 2000 \text{ filters}$
64.3	63.3	62.2	AESD: $k(x, y) = \frac{0.5}{1 + x - y } + \frac{1 + exp(-1)(xy + e)}{1 + exp(-1)(xy + e)}, \gamma = 1.0, 3000 \text{ filters}$
63.9	64.5	61.1	AESD: $k(x, y) = \frac{1+ x-y }{0.5} + \frac{1+exp(-\gamma(xy^{\intercal}+c))}{1+exp(-\gamma(xy^{\intercal}+c))}, \gamma = 1.0, 4000 \text{ filters}$

Table 3: Experimental results of various similarities. All results in above part are based on Architecture II with 1000 filters (corresponding to Row 4 in Table 2). In the bottom part, the results are based on Architecture II using the proposed metric with more filters. k(x, y) is the similarity between vector x and y. ||x|| is the L_2 norm and $||x||_1$ is the L_1 norm. xy^{T} represents the inner product of x and y. We always normalize the question and answer vectors before calculating the similarity. Highest number in each column is in **bold** font.

sources for this work are enormous. We heavily occupy a Power 7 cluster which consists of 75 machines. Each machine has 32 physical cores and each core supports 2-4 hyperthreading. The HOGWILD approach will bring some randomness due to no locking. Even with locking, the thread scheduler would alter the order of examples between runs so that randomness would still exist. Therefore, for each row in Table 2 (except for row 1 2) and Table 3, 10 experiments have been conducted on the dev set and the run with best dev score is chosen to calculate the test scores.

4. RESULTS AND DISCUSSIONS

In this section, detailed analysis on experimental results are given. From Table 2 and 3 the following conclusions can be made: (1) baseline 1 only utilizes word embeddings and base-

line 2 is based on traditional term based features. Our proposed method can reach significantly better accuracy which demonstrates the superiority of deep learning approach; (2) using separate hidden layer (HL) or CNN layers for Q and A has worse performance compared to a shared HL or CNN layer (Table 2, row 3 vs. 4, row 5 vs. 6). This is reasonable because for a shared layer network, the corresponding elements in Q and A vector are guaranteed to represent the same CNN filter convolution result while for network with separate Q and A layers, there is no such constraint and the optimizer has to learn over a set of double sized parameters. Hence the optimizer faces greater difficulty; (3) adding a HL after the CNN degrades the performance (Table 2, row 4 vs. 6 and 7). This proves that CNN already captures useful features for QA matching and unnecessary mapping the features to another space makes no sense at all; (4) increasing the CNN

filter quantity can capture more features which gives notable improvement (Table 2, row 4 vs. 8, 9 and 10); (5) two layers of CNN can represent a higher level of abstraction with wider range in the input. Hence going deeper by using two CNN layers improves the accuracy (Table 2, row 4 vs. 11); (6) effective learning in deep networks is often a difficult task. Layer-wise supervision can alleviate the problem (Table 2, row 11 vs. 12); (7) combining bigram and skip-bigram features brings gain on Test1 but not on Test2 (Table 2, row 4 vs. 13, row 8 vs. 14, row 9 vs. 15); (8) Table 3 shows that with the same model capacity, similarity metric plays an important role and the widely used cosine similarity is not the best choice for this task. The similarity in Table 3 can be categorized into three classes: L1-norm based metric which is the semantic distance of Q and A summed from each coordinate axis; L2-norm based metric which is the straight-line semantic distance of Q and A; inner product based metric which measures the angle between Q and A. We propose two new metrics that combine L2-norm and inner product by multiplication (GESD Geometric mean of Euclidean and Sigmoid Dot product) and addition (AESD Arithmetic mean of Euclidean and Sigmoid Dot product). The proposed two metrics are the best among all compared metrics. Finally, in the bottom of Table 3 it is clear that with more filters, the proposed metric can achieve even better performance.

5. RELATED WORK

Deep learning technology has been widely used in machine learning tasks, often demonstrating superior performance compared to traditional methods. Many of those applications focus on classification related tasks, e.g. on image recognition [9], on speech [10] [11] [12] and on machine translation [13] [14]. This paper is based on many prior works on utilizing deep learning for NLP tasks: Gao et al. [15] proposed a CNN based network which maps source-target document pairs to embedding vectors such that the distance between source documents and their corresponding interesting targets is minimized. Lu and Li [16] propose a CNN based deep network for a short text matching task; Hu et al. [7] also use several CNN based networks for sentence matching; Kalchbrenner et al. [17] use a CNN for sentiment prediction and question classification; Kim [18] uses a CNN in sentiment analysis; Zeng et al. [19] use a CNN for relation classification; Socher et al. [20] [21] use a recursive network for paraphrase detection and parsing; Iyyer et al. [22] propose a recursive network for factoid question answering; Weston et al. [6] use a CNN for hashtag prediction; Yu et al. [23] use a CNN for answer selection; Yin and Schütze [24] use a bi-CNN for paraphrase identification. Our work follows the spirit of many previous work in the sense that we utilize CNN to map natural language sentences into embedding vectors so that the similarity can be calculated. However this paper has conducted extensive experiments over various architectures which are not included in previous work. Furthermore, we explored different similarity metrics, skip-bigram based convolution and layerwise supervision which have not been presented in previous work.

6. CONCLUSIONS

In this paper, the spoken question answering system is studied from an answer selection perspective by employing a deep learning framework. The proposed framework does not rely on any linguistic tool and can be easily adapted to different languages or domains. Our work serves as solid evidence that deep learning based QA is an encouraging research direction. The scientific contributions can be summarized as follows: (1) creating a new QA task in the insurance domain and releasing a new corpus so that different methods can be fairly compared; (2) proposing a general deep learning framework with several variants for the QA task and comparison experiments have been conducted; (3) utilizing novel techniques that bring improvements: multi-layer CNN with layer-wise supervision, augmented CNN with discontinuous convolution and novel similarity metric that combine both L2-norm and inner product information; (4) the best scores in this paper are very promising: for this challenging task (select one answer from a pool with size 500), the top one accuracy of test corpus can reach up to 65.3%; (5) for researchers who want to proceed with this task, this paper provides valuable guidance: a shared layer structure should be adopted; no need to append a hidden layer after the CNN; two levels of CNN with layerwise training improves accuracy; discontinuous convolution sometimes can help; the similarity metric plays a crucial role and the proposed metric is preferred and finally increasing the filter quantity brings improvement.

7. REFERENCES

- [1] David A. Ferrucci, Eric W. Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya Kalyanpur, Adam Lally, J. William Murdock, Eric Nyberg, John M. Prager, Nico Schlaefer, and Christopher A. Welty, "Building watson: An overview of the deepqa project," *AI Magazine*, vol. 31, no. 3, pp. 59–79, 2010.
- [2] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in Neural Information Processing Systems* 26, pp. 3111–3119. Curran Associates, Inc., 2013.
- [3] Michael Bendersky, Donald Metzler, and W. Bruce Croft, "Learning concept importance using a weighted dependence model," in *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, New York, NY, USA, 2010, WSDM '10, pp. 31–40, ACM.
- [4] Michael Bendersky, Donald Metzler, and W. Bruce Croft, "Parameterized concept weighting in verbose queries," in *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 2011, SIGIR '11, pp. 605–614, ACM.
- [5] Yoshua Bengio, Ian J. Goodfellow, and Aaron Courville, "Deep learning," Book in preparation for MIT Press, 2015.
- [6] Jason Weston, Sumit Chopra, and Keith Adams, "#tagspace: Semantic embeddings from hashtags," in Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). 2014, pp. 1822–1827, Association for Computational Linguistics.
- [7] Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen, "Convolutional neural network architectures for matching natural language sentences," in Advances in Neural Information Processing Systems 27, Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, Eds., pp. 2042–2050. Curran Associates, Inc., 2014.
- [8] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu, "Hogwild: A lock-free approach to parallelizing stochastic gradient descent," in Advances in Neural Information Processing Systems 24, J. Shawe-Taylor, R.S. Zemel, P.L. Bartlett, F. Pereira, and K.Q. Weinberger, Eds., pp. 693–701. Curran Associates, Inc., 2011.

- [9] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," in *Intelligent Signal Processing*. 2001, pp. 306– 351, IEEE Press.
- [10] Holger Schwenk, "Continuous space language models," *Comput. Speech Lang.*, vol. 21, no. 3, pp. 492–518, July 2007.
- [11] Geoffrey Hinton, Li Deng, Dong Yu, Abdel rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara Sainath George Dahl, and Brian Kingsbury, "Deep neural networks for acoustic modeling in speech recognition," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 82–97, November 2012.
- [12] Alex Graves, Abdel-rahman Mohamed, and Geoffrey E. Hinton, "Speech recognition with deep recurrent neural networks," in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2013, Vancouver, BC, Canada, May 26-31, 2013,* 2013, pp. 6645– 6649.
- [13] Jacob Devlin, Rabih Zbib, Zhongqiang Huang, Thomas Lamar, Richard Schwartz, and John Makhoul, "Fast and robust neural network joint models for statistical machine translation," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Baltimore, Maryland, June 2014, pp. 1370–1380, Association for Computational Linguistics.
- [14] Martin Sundermeyer, Tamer Alkhouli, Joern Wuebker, and Hermann Ney, "Translation modeling with bidirectional recurrent neural networks," in *Proceedings* of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, October 2014, pp. 14–25, Association for Computational Linguistics.
- [15] Jianfeng Gao, Patrick Pantel, Michael Gamon, Xiaodong He, and Li Deng, "Modeling interestingness with deep neural networks," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, October 2014, pp. 2–13, Association for Computational Linguistics.
- [16] Zhengdong Lu and Hang Li, "A deep architecture for matching short texts," in Advances in Neural Information Processing Systems 26, C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, Eds., pp. 1367–1375. Curran Associates, Inc., 2013.
- [17] Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom, "A convolutional neural network for modelling sentences," in *Proceedings of the 52nd Annual Meeting*

of the Association for Computational Linguistics (Volume 1: Long Papers). 2014, pp. 655–665, Association for Computational Linguistics.

- [18] Yoon Kim, "Convolutional neural networks for sentence classification," in *Proceedings of the 2014 Conference* on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, October 2014, pp. 1746–1751, Association for Computational Linguistics.
- [19] Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao, "Relation classification via convolutional deep neural network," in *Proceedings of COLING* 2014, the 25th International Conference on Computational Linguistics: Technical Papers, Dublin, Ireland, August 2014, pp. 2335–2344, Dublin City University and Association for Computational Linguistics.
- [20] Richard Socher, Cliff C. Lin, Andrew Y. Ng, and Christopher D. Manning, "Parsing natural scenes and natural language with recursive neural networks," in *Proceedings of the 26th International Conference on Machine Learning (ICML)*, 2011.
- [21] Richard Socher, Eric H. Huang, Jeffrey Pennington, Andrew Y. Ng, and Christopher D. Manning, "Dynamic Pooling and Unfolding Recursive Autoencoders for Paraphrase Detection," in Advances in Neural Information Processing Systems 24. 2011.
- [22] Mohit Iyyer, Jordan Boyd-Graber, Leonardo Claudino, Richard Socher, and Hal Daumé III, "A neural network for factoid question answering over paragraphs," in *Empirical Methods in Natural Language Processing*, 2014.
- [23] Lei Yu, Karl Moritz Hermann, Phil Blunsom, and Stephen Pulman, "Deep Learning for Answer Sentence Selection," in *NIPS Deep Learning Workshop*, Dec. 2014.
- [24] Wenpeng Yin and Hinrich Schütze, "Convolutional neural network for paraphrase identification," in Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, May–June 2015, pp. 901–911, Association for Computational Linguistics.