

Exploiting Sentence Embedding for Medical Question Answering

Yu Hao,^{1*} Xien Liu,^{1*} Ji Wu,¹ Ping Lv²

¹Department of Electronic Engineering, Tsinghua University, Beijing, China
haoy15@mails.tsinghua.edu.cn, {xeliu, wuji_ee}@mail.tsinghua.edu.cn

²Tsinghua-iFlytek Joint Laboratory, iFlytek Research, Beijing, China
luping_ts@mail.tsinghua.edu.cn

Abstract

Despite the great success of word embedding, sentence embedding remains a not-well-solved problem. In this paper, we present a supervised learning framework to exploit sentence embedding for the medical question answering task. The learning framework consists of two main parts: 1) a sentence embedding producing module, and 2) a scoring module. The former is developed with contextual self-attention and multi-scale techniques to encode a sentence into an embedding tensor. This module is shortly called Contextual self-Attention Multi-scale Sentence Embedding (CAMSE). The latter employs two scoring strategies: Semantic Matching Scoring (SMS) and Semantic Association Scoring (SAS). SMS measures similarity while SAS captures association between sentence pairs: a medical question concatenated with a candidate choice, and a piece of corresponding supportive evidence. The proposed framework is examined by two Medical Question Answering (MedicalQA) datasets which are collected from real-world applications: medical exam and clinical diagnosis based on electronic medical records (EMR). The comparison results show that our proposed framework achieved significant improvements compared to competitive baseline approaches. Additionally, a series of controlled experiments are also conducted to illustrate that the multi-scale strategy and the contextual self-attention layer play important roles for producing effective sentence embedding, and the two kinds of scoring strategies are highly complementary to each other for question answering problems.

Introduction

Embedding learning in word-level has achieved much progress (Bengio et al. 2003; Mikolov et al. 2013b; Mikolov et al. 2013a; Pennington, Socher, and Manning 2014) and the pre-trained word embeddings have been almost a standard input to a certain deep learning framework for solving upstream applications, such as reading comprehension tasks (Raison et al. 2018; Wang et al. 2018; Zhang et al. 2018; Chen et al. 2017; Cheng, Dong, and Lapata 2016; Dhingra et al. 2016; Seo et al. 2016). However, learning embeddings at sentence/document level is still a very difficult task, not well solved at present. The study of sentence embedding runs along the two lines: 1) exploiting

semantic/linguistic properties (Zhu, Li, and Melo 2018; Baroni et al. 2018) obtained within sentence embeddings, and 2) designing learning methods to produce effective sentence embeddings. All of the learning methods can be generally categorized into two groups: 1) obtaining universal sentence embeddings with an unsupervised learning framework (Hill, Cho, and Korhonen 2016; Kiros et al. 2015; Le and Mikolov 2014), and 2) producing task-dependent sentence embeddings with a supervised learning framework (Palangi et al. 2016; Tan et al. 2016; Cheng, Dong, and Lapata 2016; Lin et al. 2017).

Though plenty of successful deep learning models are built at word level (word embedding), there are still some different and valuable merits obtained within sentence embedding. For example, most reading comprehension models calculate a pairwise similarity at word level to extract keywords in the answer. However, these fine-grained models maybe misled under certain circumstances, such as long paragraph with lots of single noisy words which are similar to those words that appear in the question but unrelated to the question answering. Furthermore, models built on sentence embeddings sometimes can be more interpretable. For example, we can encode a sentence into several embeddings to capture different semantic aspects of the sentence. As we known, sometimes interpretability becomes more crucial for certain real applications, such as tasks from the medical domain.

In this paper, we focus on developing supervised sentence embedding learning framework for solving medical question answering problems. To maintain model interpretation, we also adopt the self-attention structure proposed by (Lin et al. 2017) to produce sentence embeddings. The only difference is that a contextual layer is used in conjunction with the self-attention. Under certain circumstances, the valuable information resides in a unit whose size is between word and sentence. Take medical text for an instance, a large amount of the medical terminologies are entities consist of several sequential words like *Acute Upper Respiratory Infection*. It requires a flexible scale between word and sentence-level to encode such sequential words as a single unit and assign words in the unit with similar attention, instead of treating them like a bag of unrelated words, which can be misled easily by noisy words in long paragraphs when computing pairwise word similarities. For example, sentences

*These two authors contribute equally to this work.

that include *Acute Gastroenteritis* and *Acute Cholecystitis* may be considered to some extent related to question that describes *Acute Upper Respiratory Infection* because *Acute* appears in all of them, even though these sentences concentrate on totally different diseases. Therefore we propose contextual self-attention and multi-scale strategy to produce sentence embedding tensor that captures multi-scale information from the sentence. The contextual attention detects meaningful word blocks(entities) and assigns words in the same block with similar attention value. The multi-scale allows the model to directly encode sequential words as an integral unit and extracts informative entities or phrases. The contextual attention is a soft assignment of attention value, while the multi-scale strategy is a hard binding of sequential words. Even though being able to preserve more information by producing a tensor, Lin et al. simply calculate similarities between corresponding semantic subspaces and fail to capture the association between different subspaces. In an attempt to fully exploit the abundant information lies in the tensor, we propose two scoring strategies: Semantic Matching Scoring(SMS) and Semantic Association Scoring(SAS).

In the rest of this paper, we will first define the medical question answering task and introduce two datasets. Then, the supervised sentence embedding learning framework (consisting of sentence embedding producing module CAMSE and scoring module) are introduced, and a series of comparison results and some crucial analysis are presented.

MedicalQA Task Description

Here, we define the MedicalQA task with three components:

- Question: a short paragraph/document in text describing a medical problem.
- Candidate choices : multiple candidate choices are given for each question, and only one is the correct answer.
- Evidence documents: for each candidate choice, a collection of short documents/paragraphs¹ is given as evidence to support the choice as the right answer.

The goal of MedicalQA is to determine the correct answer based on corresponding evidence documents with an appropriate scoring manner.

$$(Q, \{c_1, c_2, \dots, c_{n_c}\}, \{D_1, D_2, \dots, D_{n_c}\}) \rightarrow c^*, \quad (1)$$

where n_c is the number of candidate choices for each question, c_i is the i th candidate choice, and $D_i = \{d_{i1}, d_{i2}, \dots, d_{in_e}\}$ is the set of evidence documents for the choice c_i , where $i = 1, 2, \dots, n_c$ is the index of candidate choice, and n_e is the number of evidence documents for each choice.

In the rest of this section, we will introduce two kinds of medical question answering problems, which are from real-world applications and can be transformed into MedicalQA task as defined in formula (1). The first task comes from a medical exam: the General Written Test (GWT) of National

¹In the rest of this paper, we will not specifically differentiate sentences from documents/paragraphs. These terms can be used interchangeably.

Medical Licensing Examination in China (NMLEC), and the second task is Clinical Diagnosis based on Electronic Medical Records(CD-EMR).

MedicalQA#1:NMLEC

Data source NMLEC is an annual certification exam which comprehensively evaluates doctors' medical knowledge and clinical skills. The General Written Test part of NMLEC consists of 600 multiple choice questions. Each question is given 5 candidate choices (one example is presented in Fig. 1), and only one of them is the correct answer. The exam examines more than 20 medical subjects

Question: A male patient, aged 20 years old. He had diarrhea 3 weeks ago, and 2 days later the symptoms improved but he did not mind. 1 days ago in the morning he felt weakness in limbs and pains in double legs. Gradually illness turned more serious. His family found his double eyelid cannot fully close, no dysphagia, and urinary and stool were normal. Admission examination: clear consciousness, speak normally, bilateral peripheral facial paralysis, limb muscle strength II, low muscle tension, no obvious sensory disturbance. This patient is most likely diagnosed as: ()

(A) Guillain-Barre syndrome (B) Parkinson's disease
(C) Purulent meningitis (D) Myasthenia gravis
(E) Acute myelitis

Figure 1: An example question from the General Written Test part of NMLEC.

Training/test set We collected 10 suites of the exam, totally 6,000 questions, as the test set. To avoid test questions appearing in the training set with minor variance, we dropped the training questions which are very similar to the questions from the test set, resulting in totally 250,000 medical questions as the training set. The similarity of two questions is measured by comparing Levenshtein distance(Levenshtein 1966) with a threshold ratio of 0.8.

Evidence documents The General Written Test of NMLEC mainly examines medical knowledge from medical textbooks. Therefore, we first collected totally more than 30 publications (including textbooks, guidebooks etc.) as evidence source. Then we produced evidence documents from the evidence source with a text retrieval system built upon Apache Lucene with BM25 ranking.

MedicalQA#2:CD-EMR

Data source A large amount of electronic medical records(EMRs) are collected from the outpatient department of basic level hospitals. An example of the EMRs sample is listed in Fig.2. The EMRs mainly consists of three parts:

- Chief complaint: a very brief description of the patient's illness symptom.
- History of present illness²: a more detailed interview (comprehensive description to the patient's illness) prompted by the chief complaint or presenting symptom.

²History of Present Illness, commonly abbreviated HPI, is also termed History of Presenting Complaint (HPC) in medical domain.

- Disease code: a code consisting of a combination of uppercase and numbers with a length of four. Each code indicates only one disease.

<p>Chief complaint: Sore throat for 5 days.</p> <p>History of present illness: After catching a chill 5 days ago, the patient began to suffer from sore throat, which worsened when swallowing. No nasal obstruction. No hoarseness. No cough. No fever. No expiratory dyspnea. The patient has no difficulty in somnus and eating.</p> <p>Disease code: J06.9</p>

Figure 2: An example EMR data.

Questions & candidate choices To transform the EMRs data (see Fig.2) into a standard MedicalQA task defined above, here we concatenate each chief complaint and its corresponding history of present illness, and treat the concatenation as a medical question. There are totally 98 disease codes in our data.

Training/test set We collected EMRs data during a long period from the outpatient department of basic level hospitals as the training set, and the next period data collected from the same hospitals are used as test set. Here, we use the next period data as the test set. The main purpose is to make the problem more suitable for real applications (learning from historical data, but predicting over present/future data). All training/test EMRs data are transformed into MedicalQA questions via the method mentioned above. The training set has 75265 items, and the test set has 16551 items.

Evidence documents Since MedicalQA#2 are collected from real-world EMRs data, the question description varies significantly due to the diversity of human doctors’ writing styles and deviates from textbook styles as well. Using a text retrieval system to retrieve evidence documents as done in MedicalQA#1 is not a good choice (We have a try, but the results are terrible). Here, we selected similar question samples from training set as evidence documents. All training questions are used to train a simple LSTM-MLP classifier. The output of LSTM-MLP is treated as a representation of the problem and we use it to select nearest neighbors of each problem as their supportive documents for each disease.

The Framework

According to the MedicalQA task defined above (1), the key to determining the correct answer from many candidate choices is to evaluate the supportive degree of an evidence document to the corresponding candidate choice. In this study, we consider the supportive degree into two aspects: how semantic similarity between the sentence pairs (question concatenated with a candidate choice, and the corresponding evidence document), and how association across the sentence pairs. The former is degreed with Semantic Matching Scoring (SMS) and the latter is measured with Semantic Association Scoring (SAS). In the rest of this section,

we will first introduce the sentence embedding producing module CAMSE and then present the scoring module consisting of SMS and SAS.

CAMSE

The sentence embedding producing module CAMSE is presented in Fig.3. We will introduce its details layer by layer.

Input layer & multi-scale context layer For each word of the input sentence, we lookup its embedding from pre-trained embedding matrix. Then we implement a multi-scale convolution operation on word embeddings with variable window size $i(i = 1, 2, 3, \dots, k)$. (see Fig.3 (a)).The variation of granularity enables the model to not only process single words, but also bind sequential i words as an integral representation for the potential terminology of entities or phrases, like *chronic bronchitis*. The outputs of convolution with different window sizes are processed separately with different bidirectional LSTM networks to generate semantic encodings.

Contextual self-attention layer Similar to previous self-structured sentence embedding model(Lin et al. 2017), at each scale a sentence $H^i = (h_1^i, h_2^i, h_3^i, \dots, h_n^i)$ with variable length n is encoded into a fixed-shaped 2-D feature matrix. The multi-attention mechanism attends to different semantic components of the sentence, and preserve more information than a single vector. We encode each word in the sentence into a r -dimensional attention vector, representing its significance in corresponding semantic sub-spaces.

As Fig.3 (b) shows, we first use an unbiased 1-layer feed-forward network to compress the word representation. $W_{s1}^i \in \mathbb{R}^{2u_1 \times d_a}$, d_a is the size of hidden state and u_1 is the one-direction output size of Bi-LSTM in the previous layer.

$$M^{i,1} = \tanh(H^i W_{s1}^i) \quad (2)$$

The hidden states $M^{i,1} = (m_1^{i,1}, m_2^{i,1}, \dots, m_n^{i,1})$ are processed with a 1-layer Bidirectional LSTM to integrate context information. Although the LSTM network in multi-scale context layer has contained dynamic information, we still adopt another Bi-LSTM layer here to separate the function of two Bi-LSTM networks. The first Bi-LSTM layer concentrates on semantic encoding which is further utilized when producing sentence embeddings. The output embeddings of sequential words in entities can vary significantly after the first Bi-LSTM layer, in order to preserve the diverse semantic information in a sentence. The second layer, on the other hand, focuses on detecting the meaningful word blocks and assigns words in them with similar attention values. The contextual information is incorporated so that the attention layer can better capture word blocks and treat the words in a block equally, even though their semantic embeddings might vary drastically.

$$m_t^{i,2} = \text{Bi-LSTM}(m_{t-1}^{i,2}, m_t^{i,1}) \quad (3)$$

$M^{i,2} = (m_1^{i,2}, m_2^{i,2}, \dots, m_n^{i,2})$, where $M^{i,2} \in \mathbb{R}^{n \times 2u_2}$, u_2 is the one-direction output size of Bi-LSTM in the context unit. The outputs of Bi-LSTM network at steps are then

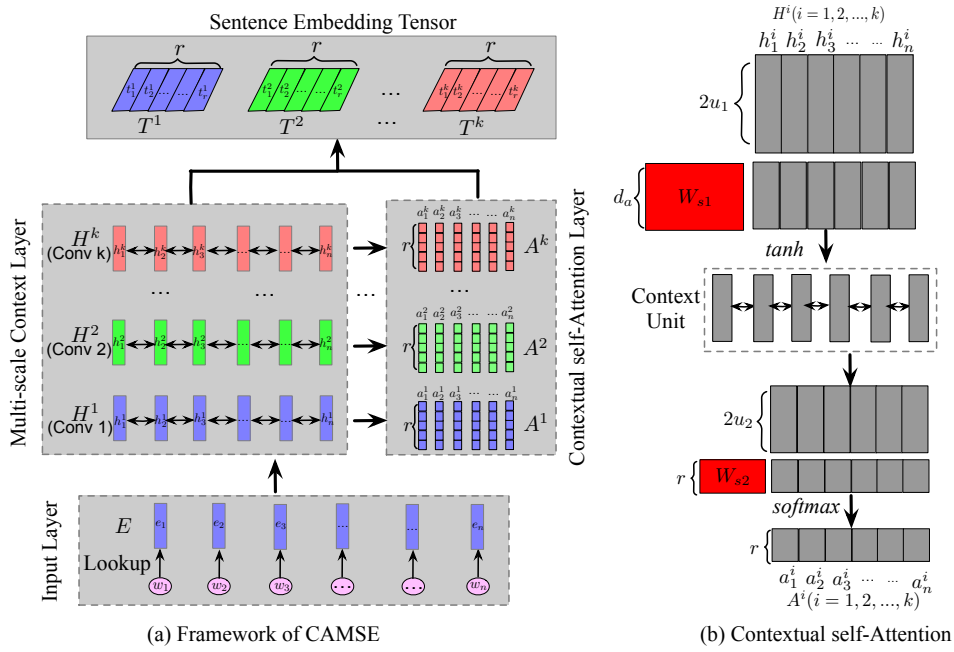


Figure 3: The framework of sentence embedding learning: CAMSE

encoded with a one-layer feed-forward network. The softmax is performed along the first dimension to ensure the attention of words in a semantic sub-space sums to 1.

$$A^i = \text{softmax}(M^{i,2}W_{s2}^i) \quad (4)$$

$W_{s2}^i \in \mathbb{R}^{2u_2 \times r}$, r is the number of semantic sub-spaces. $A^i \in \mathbb{R}^{n \times r}$ is the attention matrix. $A^i = (a_1^i, a_2^i, a_3^i, \dots, a_n^i)$, where $i = 1, 2, 3, \dots, k$, each element $a_j^i \in \mathbb{R}^{r \times 1}$ ($j = 1, 2, 3, \dots, n$) is an attention vector.

Sentence embedding tensor The attention matrix A^i is used as summation weights to summarize representations in a sentence. (see the upper part of Fig.3 (a))

$$T^i = (A^i)^T H^i \quad (5)$$

$T^i = (t_1^i, t_2^i, \dots, t_r^i)$, where $t_j^i \in \mathbb{R}^{1 \times 2u_1}$ is an embedding vector, and r is the number of semantic sub-spaces.

$T = (T^1, T^2, \dots, T^k)$ is the sentence embedding tensor generated by CAMSE (Contextual self-attention Multi-scale Sentence Embedding). The tensor $T \in \mathbb{R}^{k \times r \times 2u_1}$ automatically aligns information from the sentence in k scales and r aspects.

Scoring module

Given a pair of sentences (Q, D), we first apply CAMSE to attain sentence embedding tensors T^{s1}, T^{s2} respectively for question and document. (see Fig.4 (a)). Then a scoring function takes the tensor pair (T^{s1}, T^{s2}) as input, computing a scalar score S as supporting degree of the document.

We propose two approaches of scoring, the Semantic Matching Scoring (SMS) and Semantic Association Scoring (SAS). (see Fig. 4 (a)). These two methods can be utilized

together to boost the performance. A question-dependent gated matrix $G^{s1,i} \in \mathbb{R}^{r \times r}$ together with masks $M_{sms} \in \mathbb{R}^{r \times r}$ and $M_{sas} \in \mathbb{R}^{r \times r}$ control the information flow of two methods. As shown in Fig.4 (a). The scores of two approaches from all k scales are aggregated with a 1-layer MLP to predict a scalar score S , where $\mathbf{w}_s \in \mathbb{R}^{2k}$.

$$S = \mathbf{w}_s^T [O_{sms}^1, \dots, O_{sms}^k, O_{sas}^1, \dots, O_{sas}^k] \quad (6)$$

Here, the scalars O_{sms}^i and O_{sas}^i are the outputs of Semantic Matching Scoring and Semantic Association Scoring of the i th scale.

SMS: Semantic Matching Scoring The u th column of two embedding tensors $T_u^{s1,i}$ and $T_u^{s2,i}$, are aligned to the same semantic sub-spaces. We compute a cosine similarity for each pair of semantic sub-space embedding columns. (see Fig.4 (b))

$$S_{sms}^i(u, u) = \frac{T_u^{s1,i} \cdot T_u^{s2,i}}{\|T_u^{s1,i}\| \|T_u^{s2,i}\|} \quad (7)$$

SAS: Semantic Association Scoring As for different columns of two embedding tensors, we cannot simply compute their cosine similarities because different semantic sub-spaces are not aligned. However, we can utilize the inter-semantic relationship to exploit associations between different semantic sub-spaces. We concatenate two embeddings and send them into a 1-layer MLP to measure the correlation between these two semantics. (see Fig.4 (c)). The MLP outputs a scalar value for each semantic pair (u, v) . Different semantic pairs hold different sets of parameters $\mathbf{w}_{uv} \in \mathbb{R}^{4u_1}$.

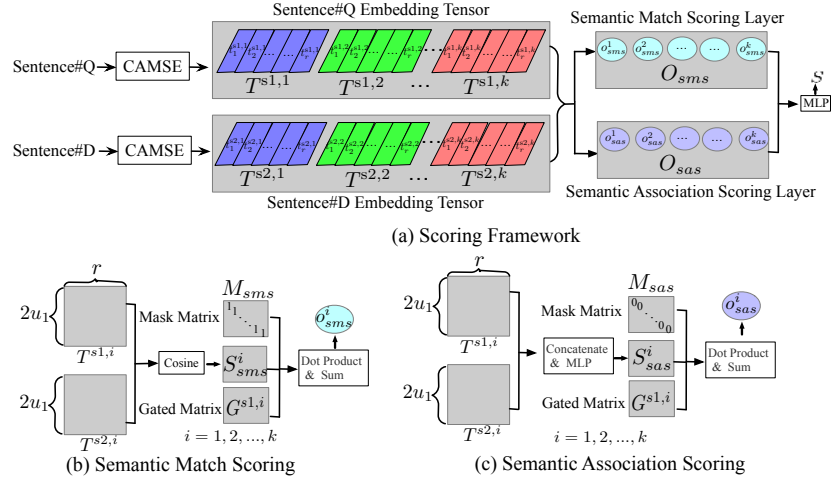


Figure 4: The framework of scoring based on sentence embedding.

$$S_{sas}^i(u, v) = \text{sigmoid}(\mathbf{w}_{uv}^T [T_u^{s1,i}, T_v^{s2,i}]) \quad (8)$$

The intuition is based on the fact that though words in different semantics are not aligned, they may frequently co-occur in the data. The fully-connected layer takes advantage of the co-occurrence, as a complement to the SMS approach. Take the clinical data for an example, descriptions for a disease contains several aspects. Some focus on particular symptoms while others narrate what triggers the disease. The cause of disease *catch a cold* has an association with symptom *cough*. The inter-semantics scoring successfully represents the association between different semantic sub-spaces(symptom and pathogen in this case).

Gated matrix

We use a matrix gate to determine which semantic pairs play pivotal roles in answer predicting, and the semantic pairs containing irrelevant information should be discarded.

$$T_{flat}^{s1,i} = \text{flatten}(T^{s1,i}) \quad (9)$$

$$G_{flat} = \text{sigmoid}(W_{g2} \tanh(W_{g1} T_{flat}^{s1,i})) \quad (10)$$

$$G = \text{reshape}(G_{flat}, [r, r]) \quad (11)$$

The mask matrices M_{sms} and M_{sas} respectively get diagonal and non-diagonal part of the matrix.

$$O_{sms}^i = \text{sum}(S_{sms}^i \odot M_{sms} \odot G^{s1,i}) \quad (12)$$

$$O_{sas}^i = \text{sum}(S_{sas}^i \odot M_{sas} \odot G^{s1,i}) \quad (13)$$

Experiments

Evaluation protocol

All the models should generate a score S for a (*Statement, Document*) pair. We sum scores of all documents belong to a candidate answer as its reliability and select the one with the highest reliability as the correct answer.

$$S_i = \sum_{e=1}^{n_e} S_{ie} \quad (14)$$

$$c^* = \arg \max_i S_i \quad (i = 1, 2, \dots, n_e) \quad (15)$$

For those sentence embedding models, we use them to generate a sentence embedding vector separately for question(statement) and the documents, which can be considered as a siamese network. We compute cosine similarity as the score for each pair. The machine comprehension models, such as R-Net, are intended for datasets like SQuAD(Rajpurkar et al. 2016) that requires an answer span in a paragraph. We modify these models by replacing the output layer with MLP layer that outputs a scalar score as the supportive degree of the document to the statement.

We report question answering accuracy on the test set. The answer predicted by model and the true answer are denoted as c_*^i and c_{true}^i for the i th question. The indicator function $\mathbb{I}(x)$ is 1 when x is True, and 0 otherwise.

$$\text{Accuracy} = \frac{\sum_{i=1}^N \mathbb{I}(c_*^i = c_{true}^i)}{N} \quad (16)$$

Pre-trained word embeddings

Word embeddings in the input layer are trained on their corresponding medical text corpus using skip-gram (Mikolov et al. 2013a). In MedicalQA#1, Word embeddings are trained on all collected medical textbooks and examination questions in train set; In MedicalQA#2, Word embeddings are trained on all collected EMRs data. The embedding's dimension is set to 200 for MedicalQA#1 and 100 for MedicalQA#2. Unseen words during testing are mapped to a zero vector.

Model settings

To save training time on the GPU, we truncate all evidence documents and questions to no more than 100 words for

MedicalQA#1 and 70 words for MedicalQA#2. For each candidate choice, only top 10 evidence documents are used to calculate the supportive score. The Bi-directional LSTM in the context layer has a dimension of 128. The size of attention encoding hidden state d_a (see Fig.3(b)) is 100. The number of semantics, r , is 15. Without any specification, in the multi-scale context layer of CAMSE framework, the size of convolution is 1,2,and 3.

Training

We put a softmax layer on top of candidate scores and use cross-entropy as loss function. Our model is implemented with Tensorflow (Abadi et al. 2016). We use Adam optimizer with exponential decay of learning rate and a dropout rate of 0.2 to reduce overfit, and the batch size is 10.

Results and analysis

We conduct a comparison of our model CAMSE with competitive baseline approaches, including sentence embedding models such as LSTM+DSSM(Palangi et al. 2015), LSTMN(Cheng, Dong, and Lapata 2016), and Self-Attention (Lin et al. 2017); and including some famous reading comprehension models, such as R-Net(Wang et al. 2017), Iterative Attention(Sordoni et al. 2016), Neural Reasoner(Peng et al. 2015), and SeaReader(Zhang et al. 2018). The comparison results over the two kinds of MedicalQA tasks are presented in Table 1. From the results, we can see that our presented model achieve remarkable gains than other sentence embedding models (LSTM+DSSM, LSTMN, Self-Attention) and is also superior to the competitive reading comprehension models (SeaReader, R-Net, Iterative Attention, and Neural Reasoner).

The performances of sentence-level models (LSTM+DSSM, LSTMN, Self-Attention) are generally poorer than word-level machine comprehension models(SeaReader, R-Net, Iterative Attention, and Neural Reasoner), indicating the difficulty of solving question answering problems with sentence embeddings. Our sentence-level approach, however, achieves even better performance compared with machine comprehension models.

Table 1: Experimental results comparison of our CAMSE model with other baseline approaches.

Models	MedicalQA#1	MedicalQA#2
LSTM-DSSM	44.1	81.5
LSTMN	45.0	81.6
Self-Attention	65.2	78.5
CAMSE	73.6	84.3
Neural Reasoner	52.5	81.1
Iterative Attention	58.7	82.1
R-Net	63.7	82.4
SeaReader	71.8	82.4

Contextual self-attention

Fig.5 shows an example of how the contextual self-attention works. The first attention attends to

information related with "neck mass"; the second attention promotes to represent question type "the most meaningful inspect for diagnosis"; while the third mainly focuses on "fine-needle aspiration cytology", an inspect method of thyroid. We also noticed that the sequential words in terminologies or phrases are equally assigned with high attention value, indicating that they are encoded as a whole unit via the contextual self-attention mechanism.

Attention 1

Male, 45 years old, have neck mass for 3 months, no discomfort, no history of tuberculosis. The mass is on the middle left outside the neck, with a size of 2.5 centimeter, active, no tenderness. Thyroid does not reach tubercle. The most meaningful inspect for diagnosis for this patient is fine-needle aspiration cytology.

Attention 2

Male, 45 years old, have neck mass for 3 months, no discomfort, no history of tuberculosis. The mass is on the middle left outside the neck, with a size of 2.5 centimeter, active, no tenderness. Thyroid does not reach tubercle. The most meaningful inspect for diagnosis for this patient is fine-needle aspiration cytology.

Attention 3

Male, 45 years old, have neck mass for 3 months, no discomfort, no history of tuberculosis. The mass is on the middle left outside the neck, with a size of 2.5 centimeter, active, no tenderness. Thyroid does not reach tubercle. The most meaningful inspect for diagnosis for this patient is fine-needle aspiration cytology.

Figure 5: An example of contextual self-attention over MedicalQA#2 dataset. Red color indicates more attention value.

Table 2: Comparison with different self-attention strategies. MASE is the framework of Multi-scale self-Attention Sentence Embedding, which lacks context unit in the contextual self-attention layer.

Attention Strategy	MedicalQA#1	MedicalQA#2
MASE	69.7	83.2
CAMSE	73.6	84.3

Table 3: Comparison with different scales in the framework CAMSE.

Multi-scale	MedicalQA#1	MedicalQA#2
Conv 1	72.1	83.9
Conv 1+2	73.1	84.1
Conv 1+2+3	73.6	84.3

Multi-scale layer

The multi-scale layer aims at discovering entities, which is medical terminologies in our case. By binding sequential words, the multi-scale mechanism is capable of generating a representation for the entity and filter out noisy single words

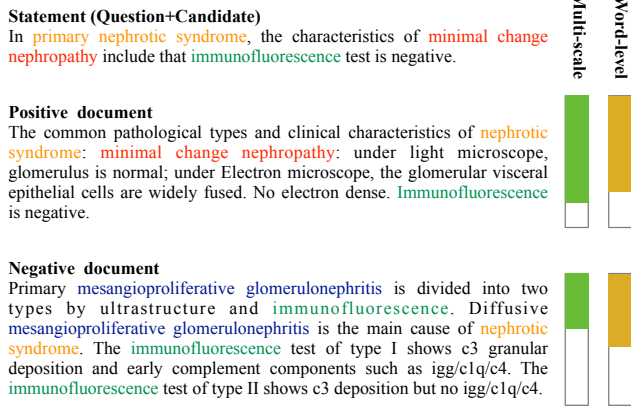


Figure 6: An example of multi-scale layer over MedicalQA#2 dataset.

that also appear in the entity. Table 3 shows the improvement of multi-scale mechanism. We observe that the improvement is greater on MedicalQA#1. The texts in MedicalQA#1 are obtained from official examination and use lots of formal terminologies, which can be well captured by multi-scale layer, while the MedicalQA#2 is subject to the diverse writing style of the doctors and harder to solve.

Fig.6 demonstrates how multi-scale mechanism outperforms the word-level models. The longer the strip is, the more support the statement receives from the document. Both the statement and positive document is about *minimal change nephropathy*, a disease of kidney, while the negative document is about *mesangioliferative glomerulonephritis*, another disease of kidney. The colored keywords in statement, including *primary nephrotic syndrome*, *immunofluorescence*, are presented in both documents. And *nephropathy* has a similar word embedding with *glomerulonephritis*. Therefore, the word-level model would encounter difficulty when trying to distinguish these sentences with similar words. The noisy words in the negative document, such as *immunofluorescence* and *glomerulonephritis*, confuse the word-level model. However, the multi-scale model is able to recognize *minimal change nephropathy* and *mesangioliferative glomerulonephritis* as integral units and easily discover that they are different diseases. The embeddings of them are distant enough so that the model can filter out irrelevant information from negative documents.

Table 4: Comparison with using different scoring strategies.

Scoring Method	MedicalQA#1	MedicalQA#2
CAMSE(SAS + SMS)	73.6	84.3
CAMSE(SAS only)	70.8	83.6
CAMSE(SMS only)	71.3	82.8

SMS and SAS The SMS measures the similarity between sentences in aligned semantic sub-space, while the SAS strategy catches association across sub-space semantics. An example from CD-EMR illustrates how they function. The

number of semantic sub-space is 5. We label the keywords of each sub-space, that is to say, words with highest attention intensity in each semantic. In parentheses are the number of sub-space the words contained in the brackets belong to.

Question : [Rhinorrhea and expectoration](5) for 4 days. [Catch a cold](2) 4 days ago and then cough, produce expectoration, have [headache](4), [sore throat](1) and [rhinorrhea](3).

Document : [Nasal obstruction, cough and expectoration](5) for 3 days. [After catching a cold](2) 4 days ago, the patient began to cough, produce expectoration and [rhinorrhea](3), with suffering from headache and [sore throat](1). The patient has [dry stool](4).

Fig.7 shows the matrix of SMS(diagonal) and SAS(non-diagonal). The SMS approach directly compare similarities, thus semantic pairs $(Q1, D1)$, $(Q2, D2)$ and $(Q3, D3)$ have higher scores, while $(Q4, D4)$ and $(Q5, D5)$ are relatively lower. The SAS approach, in this case, manages to seize the association between (sore throat, rhinorrhea), (catch a cold, rhinorrhea), (catch a cold, headache), (rhinorrhea, nasal obstruction) respectively in semantic pairs $(Q1, D3)$, $(Q2, D3)$, $(Q4, D2)$ and $(Q3, D5)$. From a large scale of data, the model discovers symptom-symptom and cause-symptom association and using it to build up connections between different aspects of the description for a disease.

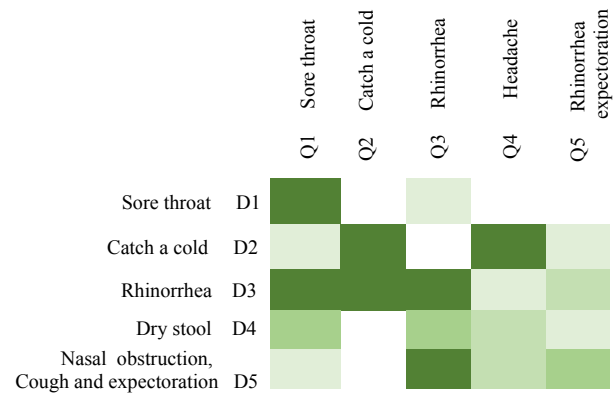


Figure 7: An example of SMS and SAS.

Conclusion

In this paper, we introduce a kind of MedicalQA task and exploit sentence embedding for this problem. A supervised learning module CAMSE is introduced to encode a sentence into an embedding tensor, and then two complementary scoring strategies Semantic Matching Scoring (SMS), and Semantic Association Scoring (SAS) are presented to exploit semantic similarity and association between a given question and the corresponding evidence document. A series of experiments are conducted on two kinds of MedicalQA datasets to illustrate that our framework can achieve significantly better performance than competitive baseline approaches. Additionally, the proposed model can maintain better model interpretation with the contextual self-attention

strategy to capture different semantic aspects at the sentence level.

Acknowledgments

We would like to thank Xiao Zhang for his help of implementing some baseline models. This work is supported by the National Key Research and Development Program of China (No.2018YFC0116800).

References

- Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin, M.; et al. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.
- Baroni, M.; Barrault, L.; Lample, G.; Kruszewski, G.; and Conneau, A. 2018. What you can cram into a single vector: Probing sentence embeddings for linguistic properties. In *ACL (1)*, 2126–2136.
- Bengio, Y.; Ducharme, R.; Vincent, P.; and Jauvin, C. 2003. A neural probabilistic language model. *Journal of machine learning research* 3(Feb):1137–1155.
- Chen, D.; Fisch, A.; Weston, J.; and Bordes, A. 2017. Reading wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, 1870–1879.
- Cheng, J.; Dong, L.; and Lapata, M. 2016. Long short-term memory-networks for machine reading.
- Dhingra, B.; Liu, H.; Yang, Z.; Cohen, W. W.; and Salakhutdinov, R. 2016. Gated-attention readers for text comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, 832–1846.
- Hill, F.; Cho, K.; and Korhonen, A. 2016. Learning distributed representations of sentences from unlabelled data. *arXiv preprint arXiv:1602.03483*.
- Kiros, R.; Zhu, Y.; Salakhutdinov, R. R.; Zemel, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, 3294–3302.
- Le, Q. V., and Mikolov, T. 2014. Distributed representations of sentences and documents. In *International Conference on Machine Learning*.
- Levenshtein, V. I. 1966. Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady* 10(1).
- Lin, Z.; Feng, M.; Santos, C. N. D.; Yu, M.; Xiang, B.; Zhou, B.; and Bengio, Y. 2017. A structured self-attentive sentence embedding.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013a. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations 2013*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119.
- Palangi, H.; Deng, L.; Shen, Y.; Gao, J.; He, X.; Chen, J.; Song, X.; and Ward, R. 2015. Deep sentence embedding using long short-term memory networks: analysis and application to information retrieval. *IEEE/ACM Transactions on Audio Speech & Language Processing* 24(4):694–707.
- Palangi, H.; Deng, L.; Shen, Y.; Gao, J.; He, X.; Chen, J.; Song, X.; and Ward, R. 2016. Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)* 24(4):694–707.
- Peng, B.; Lu, Z.; Li, H.; and Wong, K.-F. 2015. Towards neural network-based reasoning. *arXiv preprint arXiv:1508.05508*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *EMNLP*, volume 14, 1532–43.
- Raison, M.; Mazaré, P.-E.; Das, R.; and Bordes, A. 2018. Weaver: Deep co-encoding of questions and documents for machine reading. *arXiv preprint arXiv:1804.10490*.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. *ArXiv e-prints*.
- Seo, M.; Kembhavi, A.; Farhadi, A.; and Hajishirzi, H. 2016. Bidirectional attention flow for machine comprehension. In *International Conference on Learning Representations 2017*.
- Sordoni, A.; Bachman, P.; Trischler, A.; and Bengio, Y. 2016. Iterative alternating neural attention for machine reading. *arXiv preprint arXiv:1606.02245*.
- Tan, M.; Santos, C. D.; Xiang, B.; and Zhou, B. 2016. Improved representation learning for question answer matching. In *Meeting of the Association for Computational Linguistics*, 464–473.
- Wang, W.; Yang, N.; Wei, F.; Chang, B.; and Zhou, M. 2017. Gated self-matching networks for reading comprehension and question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Wang, Y.; Liu, K.; Liu, J.; He, W.; Lyu, Y.; Wu, H.; Li, S.; and Wang, H. 2018. Multi-passage machine reading comprehension with cross-passage answer verification. *arXiv preprint arXiv:1805.02220*.
- Zhang, X.; Wu, J.; He, Z.; Liu, X.; and Su, Y. 2018. Medical exam question answering with large-scale reading comprehension. *arXiv preprint arXiv:1802.10279*.
- Zhu, X.; Li, T.; and Melo, G. 2018. Exploring semantic properties of sentence embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, 632–637.