

Medical Exam Question Answering with Large-Scale Reading Comprehension

Xiao Zhang,¹ Ji Wu,¹ Zhiyang He,¹ Xien Liu,² Ying Su²

¹Department of Electronic Engineering, Tsinghua University
zhang-xi15@mails.tsinghua.edu.cn, {wuji_ee, zyhe_ts}@mail.tsinghua.edu.cn

²Tsinghua-iFlytek Joint Laboratory, Tsinghua University
{xeliu, yingsu2}@iflytek.com

Abstract

Reading and understanding text is one important component in computer aided diagnosis in clinical medicine, also being a major research problem in the field of NLP. In this work, we introduce a question-answering task called MedQA to study answering questions in clinical medicine using knowledge in a large-scale document collection. The aim of MedQA is to answer real-world questions with large-scale reading comprehension. We propose our solution SeaReader—a modular end-to-end reading comprehension model based on LSTM networks and dual-path attention architecture. The novel dual-path attention models information flow from two perspectives and has the ability to simultaneously read individual documents and integrate information across multiple documents. In experiments our SeaReader achieved a large increase in accuracy on MedQA over competing models. Additionally, we develop a series of novel techniques to demonstrate the interpretation of the question answering process in SeaReader.

1 Introduction

Natural language understanding is a crucial component in computer aided diagnosis in clinical medicine. Ideally, a computer model could read text to communicate with human doctors and utilize knowledge in text materials. Recent advances in deep-learning-inspired approaches have taken the state-of-the-art of various NLP tasks to a new level. However, directly reading and understanding text is still a challenging problem, especially in complex real-world scenarios.

Reading comprehension is a task designed for reading text and answering questions about it. An understanding of natural language and basic reasoning is required to tackle the task. The task of reading comprehension has been gaining rapid progress with the proposal of various datasets such as SQuAD (Rajpurkar et al. 2016) and many successful models.

Real-world scenarios for reading comprehension are usually much more complex. Unlike in most datasets, one does not have paragraphs of text already labeled as containing the answer to the question. Rather, one needs to find and extract relevant information from possibly large-scale text ma-

terials. Previous work (Miller et al. 2016) and (Chen et al. 2017) reads from Wikipedia dump to answer questions created from knowledge base entries. Another major challenge in real-world scenarios is that questions are often harder and more versatile, and the answer is less likely to be directly found in text.

We propose the MedQA, our reading comprehension task on clinical medicine aiming at simulating a real-world scenario. Computers need to read from various sources in aided diagnosis, such as patients' medical records and examination reports. They also need to exploit knowledge found in text materials such as textbooks and research articles. We assembled a large collection of text materials in MedQA as a source of information, to learn to read large-scale text.

Questions are taken from medical certification exams, where human doctors are evaluated on their professional knowledge and ability to make diagnosis. Questions in these exams are versatile and generally requires an understanding of related medical concepts to answer. A machine learning model must learn to find relevant information from the document collection, reason over them and make decisions about the answer.

We then propose our solution SeaReader: a large-scale reading comprehension model based on LSTM network and dual-path attention architecture. The model addresses challenges of the MedQA task from two main aspects: 1) **Leveraging information in large-scale text**: we propose a dual-path attention architecture which uses two separate attention paths to extract relevant information from individual documents as well as compare and extract information across multiple documents. Extracted information is reasoned over and integrated together to determine the answer. 2) **End-to-end training on weak labels**: although we only have labels with very little information, we still managed to train our model end-to-end. Our SeaReader features a modular design, with a clear interpretation of different levels of language understanding. We also propose a series of novel techniques that we call Interpretable Reasoning. This increases the interpretability of complex NLP models like our SeaReader.

2 Related Work

Our work is closely related to question answering and reading comprehension in the field of NLP.

Table 1: Questions by category

Category	Description	Ratio	Example
A1	Single statement best choice: single statement questions, 1 best answer, 4 incorrect or partially correct answers	36.8%	Question: The pathology in the pancreas of patients with type 1 diabetes is: Candidate Answers: a. Islet cell hyperplasia b. Islet cell necrosis c. interstitial calcification d. Interstitial fibrosis e. Islet cell vacuolar degeneration
B1	Best compatible choice: similar to A1, with a group of candidate answers shared in multiple questions	11.2%	
A2	Case summary best choice: questions accompanied by a brief summary of patient’s medical record, 1 best choice among 5 candidate answers	35.3%	Question: Male, 24 years old. Frontal edema, hematuria with cough, sputum with blood for 1 week. bp 160/100 mmhg, urinary protein (++) , rbc 30/ hp, ... (omitted text) Ultrasound show kidney size increase. At present the most critical treatment is: Candidate Answers: a. hemodialysis b. prednisone c. plasma exchange d. gamma globulin e. prednisone combined with cyclophosphamide
A3/A4	Case group best choice: similar to A2, with information shared among multiple questions	16.7%	

Reading Comprehension Tasks Datasets are crucial for developing data-driven reading comprehension models. The Text REtrieval Conference (TREC) (Voorhees and Tice 2000) evaluation tasks and the MCTest (Richardson, Burges, and Renshaw 2013) are successful early efforts at creating datasets for answering questions based text documents. Large-scale datasets have become more popular in the era of deep learning. The CNN/DailyMail dataset (Hermann et al. 2015) creates a combination of over 1 million questions on news articles by removing words in summary points. The Children’s Book dataset takes a similar approach and use excerpts from children’s books as reading materials. The SQuAD (Rajpurkar et al. 2016) dataset features 100K Wikipedia articles and human created questions. Answers are spans of text instead of single words in the reading passage, which creates more challenge in answer selection.

Some larger-scale question answering datasets aim to create a more real-world situation. For example, the Wiki-Movies benchmark (Miller et al. 2016) asks questions about movies, where Wikipedia articles and knowledge-base (OMDb) can be leveraged to answer the question. The webQA dataset (Li et al. 2016) collects 42,000 questions about daily life asked by users on QA site. The MS MARCO dataset (Nguyen et al. 2016) takes query logs of the users of Bing search engine as questions and uses passages extracted from web documents as reading material.

Neural Network Models for Reading Comprehension Numerous models have been proposed for this kind of tasks. Earlier attempts directly use LSTM to process text and uses attention to fetch information from the passage representation (Attentive Reader and the Impatient Reader, Hermann et al.2015). Later models largely rely on LSTMs and attention as the main building blocks. The Attention Sum Reader (Kadlec et al. 2016) predicts candidate word by summing attention weights on occurrences of the same word. The Gated Attention Reader (Dhingra et al. 2016) uses gating units to combine attention over the query into the passage. It also performs multi-pass readings over the document.

The Match-LSTM (Dhingra et al. 2016) uses separate LSTM layers to preprocess text input and predict answer. A span is selected from passage by predicting the start and end positions of the answer. EpiReader (Trischler et al. 2016)

differs from the above approaches by factoring question-answering into a two-stage process of candidate extraction and reasoning. Convolutional neural networks are used to generate sentence encodings in the reasoning of hypotheses.

Several state-of-the-art models make further improvements and are competitive on the challenging SQuAD dataset. BiDAF (Seo et al. 2016) incorporates attention from passage to query as well as from query to passage. R-NET (Wang et al. 2017) adds a passage-to-passage self attention layer on top of question-passage attention. The AoA Reader (Cui et al. 2016) shares the idea of utilizing both passage to query and query to passage attention, multiplying them together to make final prediction.

However, these models largely restrict themselves to selecting word(s) from the given passage, by modeling at the resolution of words. This makes them less straightforward to be applied to other tasks involving reading and understanding. They also rely on pre-selected relevant documents—a requirement not easily met in real-world settings.

3 The MedQA Task

The MedQA task answers questions in the field of clinical medicine. Rather than relying on memory and experience as human doctors do, a machine learning model may make use of a large collection of text materials to help find supporting information and reason about the answer.

Task Definition

The task is defined by its three components:

- **Question:** questions in text, possibly with a short description of the situation (of the patient).
- **Candidate answers:** multiple candidate answers are given for each question, of which one should be chosen as the most appropriate.
- **Document collection:** a collection of text material extracted from a variety of sources organized into paragraphs and annotated with metadata.

The goal is to determine the best candidate answer to the question with access to the documents.

Data

Question and Answers The goal of MedQA is to ask questions relevant in real-world clinical medicine that requires the ability of medical professionals to answer. We used the National Medical Licensing Examination¹ in China as a source of questions. The exam is a comprehensive evaluation of professional skills of medical practitioners. Medical students are required to pass the exam to be certified as a physician in China.

The General Written Test part of the exam consists of 600 multiple choice problems over 5 categories (see Table 1). We collected over 270,000 test problems from the internet and published materials such as exercise books. The problems do not have to have appeared in past exams. We filtered out any incomplete or duplicate problems. The statistics of the final problem set are shown in Table 2.

For training/test split, we created a test set as similar as possible to past exam problems, to approximate evaluation in real exams. A small subset of problems was chosen based on the source and context in which they appear. This might indicate their possible appearance in past exams, or they might be closely related to past exam problems. These problems are further split into valid/test sets. The remaining problems that are similar to any one problem in the valid/test set are removed to ensure that one cannot solve test problems by merely remembering problems in the training set. The similarity of two problems is measured by comparing Levenshtein distance (Levenshtein 1966) of questions with a threshold ratio of 0.9.

Table 2: Data statistics

	training	valid	test
Number of problems	222,323	6,446	6,405
<i>Problems</i>			
Average length of question (words)			27.4
Average length of answer (words)			4.2
Candidate answers per problem			5
<i>Documents</i>			
Number of documents			243,712
Average length of document (words)			43.2
Average number of tags in document metadata			3.8

Documents Human knowledge in almost every modern profession is extensively encoded in text. Medical students obtain their knowledge from a volume of textbooks during years of training. A machine learning model, however, can easily access a large collection of text materials at their disposal. We prepared text materials from a total of 32 publications including textbooks, reference books, guidebooks, books for text preparation, etc. These books cover a wide range of topics in clinical medicine.

Text material is extracted from these books, divided by paragraph, and annotated with metadata. Metadata include

¹<http://www.nmec.org.cn/EnglishEdition.html>

Text

In the fetal period, the development of the nervous system is ahead of other systems. Neonatal brain weight already reaches about 25% of adult brain weight, and the number of nerve cells is close to that of adults, but the dendrites and axons are less and shorter. The increase in brain weight after birth is mainly due to the increase in neuronal volume, the increase and lengthening of dendrites, and the formation and development of nerve myelin.

Metadata

growth and development, neuropsychological development, nervous system development

Figure 1: An example document

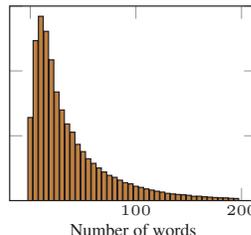


Figure 2: Length distribution of documents

contextual information like tags, book titles and chapter titles. An example document is shown in Figure 1. Basic statistics of documents after pre-processing are given in Table 2 and Figure 2.

Task Analysis

The MedQA task poses unique challenges for language understanding, especially compared with existing reading comprehension datasets. We identify several major challenges of MedQA as below:

- **Professional knowledge** Unlike most reading comprehension tasks—where question answering rely largely on commonsense reasoning and basic understanding of language—MedQA asks questions in a field of sophistication, that usually requires a thorough understanding of the field for human to answer.
- **Diversity of questions** The field of clinical medicine is diverse enough for questions to be asked about a wide range of topics. Questions can also be asked in various facets, for example: given *a description of a patient's condition*, a question might ask for *the most probable diagnosis / the most appropriate treatment / the examination needed / the mechanism of certain condition*, etc.
- **Determining the best answer** Choosing the best answer means that the unselected answers are *not* necessarily incorrect. The model must learn to evaluate individual answers and learn to make comparison.
- **Reading large-scale text** Retrieving relevant information from large-scale text is more challenging than reading a short piece of text. Furthermore, in MedQA passages from textbooks often do not directly give answers to questions, especially for *case* problems. One must discern relevant information scattered in passages, and determine the relevance of each piece of text.

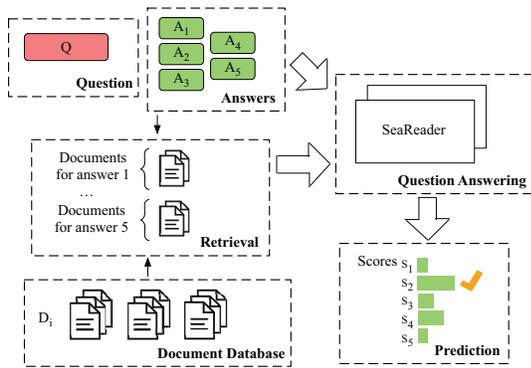


Figure 3: System overview

- Reasoning over facts** Reasoning is often required to answer the questions. This includes natural language reasoning where one recognizes lexical or syntactic variations and reasoning over facts to decide whether given document(s) support an answer to the question. Below is an example of a problem requiring reasoning over multiple facts, like those in (Weston et al. 2015):

Question: The most effective treatment for a 70 years old patient with Parkinson’s disease is:

Document 1: Commonly used drugs are anticholinergic drugs phenanthrene, amantadine, levodopa and compound levodopa

Document 2: Phenanthrene is mainly suitable for those with obvious tremble, but is more used in young patients, elderly patients should be used with caution

4 The SeaReader Framework

Our proposed solution to the MedQA task includes a document retrieval system and a neural network based question answering model. Text retrieval is used to narrow down possibly related documents that are then fed into the SeaReader where the reasoning and decision-making occur. See Figure 3 for an overview of the framework.

Document Retrieval

Given a problem, we select a subset of documents that are likely to be relevant using a text retrieval system. The text retrieval system is built upon Apache Lucene, using inverted index lookup (metadata included in documents) followed by BM25 ranking. For each problem, we perform retrieval for each answer individually by pairing it with the question. We take the intersection of the retrieved document for question and answer and keep the top-N documents. To account for the different nature of various text sources, we try to select an equal number of documents from each type of books. We found that this promotes diversity in the selected documents and improves the final performance of question answering.

The SeaReader Model

The Support Evidence Analysis Reader (SeaReader) model is designed to find relevant evidence from documents given a question and a candidate answer and perform analysis and reasoning to determine the correctness of the answer.

SeaReader uses LSTM networks to model context representations of text. Attention is used extensively in the dual-path attention architecture to model information flow between question and documents, and across multiple documents. Information from multiple documents is fused together to make the final prediction. SeaReader also has a modular design that facilitates interpretation of the reasoning process, as discussed in the next section.

Input layer The input to SeaReader is a question-answer pair (Q, A) (the answer being one of the candidate answers given in the problem) and the set of documents $\{D_0, D_1, \dots, D_N\}$ returned by the retrieval system. The answer is appended to the question to form a statement S . After word-embedding lookup, we have tensors representing statement $S \in \mathbb{R}^{L_Q \times d}$ and documents $D \in \mathbb{R}^{N \times L_D \times d}$. L_Q and L_D are the maximum length of S and D_i respectively, and d is the dimension of word-embedding.

Context layer The word-level representations are then processed by a bi-directional LSTM layer to model contextual representations. Leaving out this layer, replacing it with a convolutional layer or GRU (Cho et al. 2014) all results in performance drop, indicating the importance of long-range context in representing semantics.

Dual-path attention layer We would like the model to learn to identify and extract relevant information from long documents. The dual-path attention architecture allows for extracting information from two perspectives: in the question-centric (Q-centric) path, information related to the question is extracted from documents and is aligned with the question. Information from each document is processed individually. In the document-centric (D-centric) path, we take the perspective of documents: information from the question is integrated to documents, then information from other documents is compared and integrated. Interaction of supporting facts in multiple documents is captured in this path.

Similar to several state-of-the-art works (Cui et al. 2016; Seo et al. 2016; Xiong, Zhong, and Socher 2016), we start by computing a matching matrix as the dot-product of the context embeddings of the question and every document:

$$M_n(i, j) = S(i) \odot D_n(j) \quad (1)$$

In the question-centric path, attention is performed column-wise on the matching matrix. Every word $S(i)$ in question-answer gets a summarization read $R_n^Q(i)$ of related information in document:

$$\alpha_n(i, j) = \text{softmax}(M_n(i, 1), \dots, M_n(i, L_D))(j) \quad (2)$$

$$R_n^Q(i) = \sum_{j=1}^{L_D} \alpha_n(i, j) D_n(j) \quad (3)$$

In the document-centric path, row-wise attention is performed to read related information in the question. Next, cross-document attention is performed on attention reads of all the documents (\oplus represents vector concatenation):

$$M'_{mn}(i, j) = (D_m(i) \oplus R_m^D(i)) \odot (D_n(j) \oplus R_n^D(j)) \quad (4)$$

$$\beta_{mn}(i, j) = \text{softmax}(M'_{m1}(i, 1), \dots, M'_{m1}(i, L_D), \dots, M'_{mN}(i, 1), \dots, M'_{mN}(i, L_D))(j) \quad (5)$$

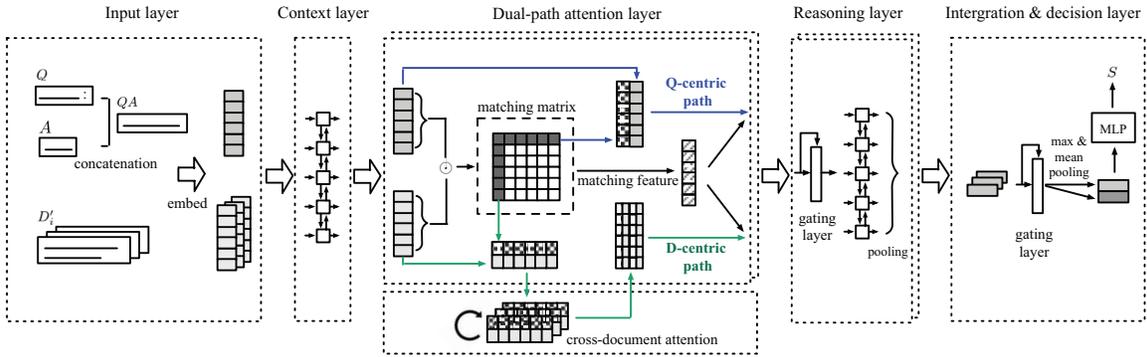


Figure 4: The SeaReader model

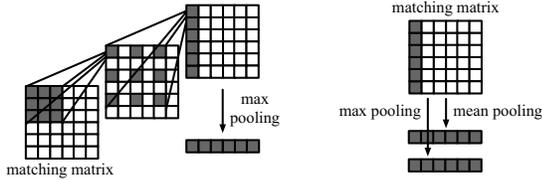


Figure 5: Extract matching features. Left: CNN extractor, right: pooling extractor

$$R_m^{D'}(i) = \sum_{n=1}^N \sum_{j=1}^{L_D} \beta_{mn}(i, j) (D_n(j) \oplus R_n^D(j)) \quad (6)$$

The cross-document attention extracts relevant information from other documents based on the current document, which enables integration of information from multiple documents.

We introduce matching features F as a complement to attention reads. Softmax in attention destroys absolute matching strength, e.g. “clinic” should matches “hospital” better than “physician”, regardless of the accompanying words in the document. We thus extract matching features directly from the matching matrix, using a two-layer convolutional network on M (see Figure 5 left). Max-pooling is performed between and after convolution layers. The second layer used dilated convolution to keep resolution at the word-level.

The convolution captures patterns in the matching matrix at a significant cost of computation. We also experimented a simpler design with similar performance gain: extracting matching feature by max-pooling and mean-pooling rows and columns in M (Figure 5 right).

Reasoning layer The reasoning layer takes attention reads from the Q-centric and the D-centric path as well as matching features. It then uses a bi-directional LSTM layer to reason on the question/document level. A gating layer is applied before LSTM to decide whether a word should contribute to reasoning. The gate value is computed from contextual embedding, and is multiplied to all input features.

The outputs of LSTM represent the support of the document to the statement, which is summarized into a single vector by a max-pooling over the sequence.

Integration & decision layer To integrate support from multiple documents, feature vectors first go through a gating layer similar to that in the reasoning layer. Gating decides relevant documents and suppresses irrelevant ones. Next, max-pooling and mean-pooling are performed to further summarize the support of all the documents. The intuition of using two different pooling together is based on the assumption that best candidates should have *better* as well as *more* support, respectively. A multi-layer feed-forward network is used to predict a scalar score. The candidate answer having the largest score value is chosen as the best answer predicted by the model.

Interpretable Reasoning

Neural network models are sometimes regarded as black-box optimization for difficulty interpreting model’s behavior. Interpretability is recognized as invaluable for analyzing and improving the model. Here we present a series of novel techniques to improve the interpretability of our SeaReader, which can also be applied to general neural network based NLP models.

Gating with importance penalty The value of a scalar gate can usually be interpreted as the importance of the gated information. However, putting a gate somewhere in the model does not always make its output meaningful. The gate can be passed-through if it does not help optimization. We introduce a regularization term in the objective function to use a gate for inspection purposes:

$$loss = loss_{task} + c \cdot \max\left(\frac{1}{L} \sum_i^L g(i) - t, 0\right) \quad (7)$$

The added term restricts the average gate value to be below a threshold of t (set to 0.7 in our experiments). The model must now learn to suppress unimportant information by giving lower gate values to those vectors.

Noisy gating It is sometimes desirable for a model to rely more on strong evidence than weak evidence because the latter contains more noise and makes the model more prone to overfitting. Gating helps to differentiate strong evidence from weaker evidence, and the effect can be reinforced by adding Gaussian noise after gating:

$$X_{out} = gate(X_{in})X_{in} + \sigma(0, s) \quad (8)$$

A high gate value helps a strong evidence stand out from noise. Weaker evidence become harder to utilize in the presence of noise. We found strong evidence getting higher gate values and otherwise for weak evidences than without the added noise. This helps interpreting gate values and can improve generalization performance.

L_{21} regularized embedding learning Word embeddings often have the largest number of parameters in a NLP model. When there is little labeled information for a complex task, it is difficult to learn embeddings well. In our experiments on MedQA, learning word embeddings directly suffers from severe overfit, even with pre-trained embedding as initialization. Fixed word embedding results in decent performance, but leads to an underfit model. To address this problem, we introduce a delta embedding term and adapt L_{21} regularization that is often used in structural optimization (Zhou, Chen, and Ye 2011; Kong, Ding, and Huang 2011):

$$\mathbf{w} = \mathbf{w}_{\text{skip-gram}} + \mathbf{w}_{\Delta} \quad (9)$$

$$\text{loss} = \text{loss}_{\text{task}} + c \sum_{i=1}^n \left(\sum_{j=1}^d w_{\Delta ij}^2 \right)^{\frac{1}{2}} \quad (10)$$

Intuitively, the model learns to refine skip-gram embeddings for supervised task while avoiding overfit by modifying only a few words for as little as possible. This does not only improve performance, but also lets us interpret semantics learned from the task, by inspecting *which* words are adjusted and *how* they are adjusted.

5 Experiments

Experimental Setup

Word embedding Word embedding is trained on a corpus combining all text from the problems in the training set and the document collection using skip-gram (Mikolov et al. 2013). The dimension is set to 200. Unseen words during testing are mapped to a zero vector.

Model settings All documents are truncated to no more than 100 words before processed by SeaReader. Although leading to a minor performance drop, this greatly accelerates the experiments by saving training time on the GPU. In most experiments, we retain 10 documents for each candidate answer, a total of 50 documents per problem.

Bidirectional LSTMs in the context layer and the reasoning layer all have a dimension of 128. Parameters are shared between LSTMs processing question and documents in the context layer. A single layer feed-forward network is used in the decision layer because more layers did not further improve performance.

Training We put a softmax layer on top of candidate score predictions and use cross entropy w.r.t. the groundtruth choice as objective function. Our model is implemented using Tensorflow (Abadi et al. 2016). Adam optimizer is used with $\epsilon = 10^{-6}$ to stabilize training. Exponential decay of learning rate and dropout rate of 0.2 was used to reduce overfit. We used a batch size of 15, which already contains 750 documents per batch and is the maximum allowed to train on a single GPU.

Baseline Approaches

To compare the performance of our SeaReader with existing reading comprehension models, we selected a few models with different considerations and adapted them to our MedQA task:

- **Iterative Attention** (Sordoni et al. 2016): This is one of the rare models not extensively tailored for cloze (or span) style tasks. It uses attention to read from question and document alternatively, and the read is performed iteratively with the final state used to make prediction. The model is directly applicable to our MedQA task.
- **Neural Reasoner** (Peng et al. 2015): A framework for reasoning over natural language sentences, using a deep stacked architecture. It can extract complex relations from multiple facts to answer questions. The model is a natural fit to our task and is straightforward to apply.
- **R-NET** (Wang et al. 2017): A recent reading comprehension model, achieving state-of-the-art single model result on the challenging SQuAD dataset². It stacks a question-to-document attention layer and a document-to-document attention layer. We replaced the final prediction layer with a pooling layer to generate a scalar score.

Experimental Results

Table 3: Results (accuracy) of SeaReader and other approaches on MedQA task

	Valid set	Test set
Iterative Attention	60.7	59.3
Neural Reasoner	54.8	53.0
R-NET	65.2	64.5
SeaReader	73.9	73.6
SeaReader (ensemble)	75.8	75.3
Human passing score ³	60.0 (360/600)	

Quantitative Results We evaluated model performance by accuracy of choosing the best candidate. The results are shown in Table 3. Our SeaReader clearly outperforms baseline models by a large margin. We also include the human passing score as a reference. As MedQA is not a common-sense question answering task, human performance relies heavily on individual skill and knowledge. A passing score is the minimum score required to get certified as a Licensed Physician and should reflect a decent expertise in medicine.

Test performance is broken down by problem category in Figure 6. Models with extensive word-level attention (SeaReader and R-NET) win on statement type problems (A1, B1), indicating better ability at capturing details and fine-grained reasoning. The dual-path attention architecture in SeaReader gives a major boost in performance, especially

²<https://rajpurkar.github.io/SQuAD-explorer/>, by the time of writing this paper

³of the year 2016

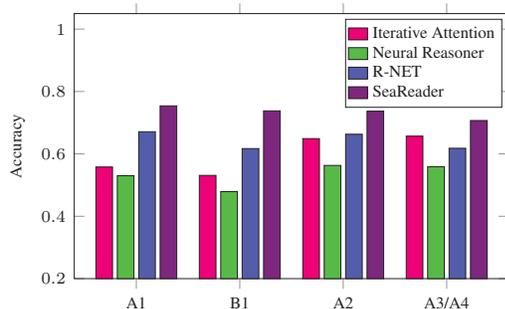


Figure 6: Accuracy by problem category

at solving more complex problems (B1, A3/A4, where information is mixed/shared among problems), showing the effectiveness of multi-perspective information extraction and integration of information in multiple documents.

Table 4: Test performance with different number of documents given per candidate answer

Number of documents	top-1	top-5	top-10	top-20
SeaReader accuracy	57.8	71.7	73.6	74.4
Relevant document ratio	0.90	0.54	0.46	0.29

Test performance clearly increases as more documents is given to SeaReader as input (see Table 4). To discover the general relevancy of documents returned by our retrieval system, we hand-labeled the relevancy of 1000 retrieved documents for 100 problems. The ratio of documents containing relevant information in the top-N documents is given in Table 4. We notice that there is still performance gain using as many as 20 documents per candidate answer, while the ratio of relevant documents is already low. This illustrates SeaReader’s ability to discern useful information from large-scale text and integrate them in reasoning and decision-making.

Discussion In the extensive architecture search developing SeaReader on MedQA task, we discovered that: 1) MedQA is a complex task with weak labels, and the best model design in such a scenario is a modular architecture without excessive depth. Our SeaReader is not only the best performing but also the fastest to converge in training. Multi-task learning is another important direction in such a scenario. 2) Extensive use of attention compensates for lack of depth, and intra-document and inter-document attention helps to model information flow at different levels, which is crucial for utilizing large-scale text inputs.

Interpret Reasoning The modular design combined with Interpretable Reasoning extensions enabled convenient interpretation of the decision-making process in SeaReader. Using noisy gating on document gating layer, the contribution of each document is clearly represented by gate value (see Figure 7 for an example). Note that the most contributing document is a semantically relevant one, without sharing many common words with question and answer.

Question: Male, 40 years old, gross hematuria with renal colic, ultrasound found stone in right kidney with size 0.6 cm and smooth surface, mild hydronephrosis. The treatment should be:
Answer: Non-surgical treatment
Documents:

- Complications: After lithotripsy the majority of patients has transient gross hematuria, which generally does not need special treatment. Renal hematoma formation is relatively rare and can be treated without surgery.
- One of the characteristics of hematuria caused by kidney tuberculosis is that it frequently happens after a period of bladder irritation, and is more common to have terminal hematuria, which is different from hematuria caused by other urinary diseases.
- Renal postoperative oliguria or no urine: they are all urological diseases, and can be diagnosed by medical history, physical examination, urology and imaging examinations. It can be clarified whether bilateral hydronephrosis is caused by bladder urinary retention or bilateral ureteral obstruction.
- Smooth stones <0.6 cm with no urinary tract obstruction or infection, also being pure uric acid stones or cystine stones, can first be treated with non-surgical therapy. 90% of smooth stones with diameter <0.4 cm can self-discharge.
- Stones <0.6 cm, smooth, no obstruction and infection, pure uric acid stones or cystine stones, can be treated by with medication to discharge or dissolve the stone.
- Focal clearance: suitable for renal tuberculosis abscess without connection to the renal pelvis. For those treated with anti-TB drugs for 3 to 6 months without effect, surgical removal of renal tuberculosis lesions can be applied.

Figure 7: Document-level gating visualization. Bar length is proportional to gate value. Documents are truncated for better visualization.

Statement1:
 Female , 23 years old . Previously healthy ; sudden hemoptysis about 500 ml . Examination: no heart or lung abnormalities ; chest x-ray film show lung lower texture thickening . The diagnosis should be bronchodilator

Statement2:
 7-month-old baby, born in hospital ; according to his/her immune characteristics and the occurrence of infectious diseases, the planned immunization and vaccination should include hepatitis B vaccine

Figure 8: Word-level gating visualization. A darker color indicates larger gate value

The word gating layer gives relative importance of each word when regulated with importance penalty (see Figure 8). Same words can get different importance by the context they appear in.

Regulating embedding with L_{21} loss allowed us to discover the “mostly learned”, or most relevant words for the task. Figure 9 shows that the most significant words for MedQA are words representing logical relations plus some commonly used terms in diagnosis.

Attention weights in SeaReader provide direct illustration about matching and relevance between words in contexts. One is able to examine the extracted information from question, documents and across documents. We omit detailed example here for brevity of the paper.

6 Conclusion

In this work, we introduced a question-answering task MedQA on clinical medicine and our reading comprehension model SeaReader. The unique challenge of MedQA is to solve advanced exam problems with large-scale text reading. Our solution SeaReader uses a dual-path attention architecture to integrate information from multiple text documents. Experiments validated its effectiveness especially at solving complex problems and utilizing information. We also propose techniques to help interpret question-answering models, which help to explain the reasoning process of SeaReader. We hope this work contributes to real-world

⁴these are translations of corresponding Chinese words

the above	not
except	diagnosis
not	acute
best	normal
treatment	not
identify	early stage
chronic	besides
small	prefererd
exam	high
obvious	large amount

Figure 9: Top-20 words learned with delta embedding⁴. Line length is proportional to norm of w_{Δ}

question answering and the application of deep-learning approaches in medicine.

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*.
- 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.
- Cho, K.; van Merriënboer, B.; Bahdanau, D.; and Bengio, Y. 2014. *On the properties of neural machine translation: Encoder-decoder approaches*.
- Cui, Y.; Chen, Z.; Wei, S.; Wang, S.; Liu, T.; and Hu, G. 2016. Attention-over-attention neural networks for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, 593–602.
- 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*, 1832–1846.
- Hermann, K. M.; Kocisky, T.; Grefenstette, E.; Espeholt, L.; Kay, W.; Suleyman, M.; and Blunson, P. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, 1693–1701.
- Kadlec, R.; Schmid, M.; Bajgar, O.; and Kleindienst, J. 2016. Text understanding with the attention sum reader network. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 908–918.
- Kong, D.; Ding, C.; and Huang, H. 2011. Robust nonnegative matrix factorization using l21-norm. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, 673–682. ACM.
- Levenshtein, V. I. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, 707–710.
- Li, P.; Li, W.; He, Z.; Wang, X.; Cao, Y.; Zhou, J.; and Xu, W. 2016. Dataset and neural recurrent sequence labeling model for open-domain factoid question answering. *arXiv preprint arXiv:1607.06275*.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations 2013*.
- Miller, A.; Fisch, A.; Dodge, J.; Karimi, A.-H.; Bordes, A.; and Weston, J. 2016. Key-value memory networks for directly reading documents. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 1400–1409.
- Nguyen, T.; Rosenberg, M.; Song, X.; Gao, J.; Tiwary, S.; Majumder, R.; and Deng, L. 2016. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*.
- Peng, B.; Lu, Z.; Li, H.; and Wong, K.-F. 2015. Towards neural network-based reasoning. *arXiv preprint arXiv:1508.05508*.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Richardson, M.; Burges, C. J.; and Renshaw, E. 2013. Mctest: A challenge dataset for the open-domain machine comprehension of text. In *EMNLP*, volume 3, 4.
- 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*.
- Trischler, A.; Ye, Z.; Yuan, X.; and Suleman, K. 2016. Natural language comprehension with the epireader. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 128–137.
- Voorhees, E. M., and Tice, D. 2000. Overview of the trec-9 question answering track. In *TREC*.
- 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*.
- Weston, J.; Bordes, A.; Chopra, S.; Rush, A. M.; van Merriënboer, B.; Joulin, A.; and Mikolov, T. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. In *International Conference on Learning Representations 2016*.
- Xiong, C.; Zhong, V.; and Socher, R. 2016. Dynamic coattention networks for question answering. In *International Conference on Learning Representations 2017*.
- Zhou, J.; Chen, J.; and Ye, J. 2011. Malsar: Multi-task learning via structural regularization. *Arizona State University* 21.