

# JEC-QA: A Legal-Domain Question Answering Dataset

Haoxi Zhong,<sup>\*1</sup> Chaojun Xiao,<sup>\*1</sup> Cunchao Tu,<sup>1</sup> Tianyang Zhang,<sup>2</sup> Zhiyuan Liu,<sup>†1</sup> Maosong Sun<sup>1</sup>

<sup>1</sup>Department of Computer Science and Technology  
Institute for Artificial Intelligence, Tsinghua University, Beijing, China  
Beijing National Research Center for Information Science and Technology, China

<sup>2</sup>Beijing Powerlaw Intelligent Technology Co., Ltd., China  
zhonghaoxi@yeah.net, {xcjthu, tucunchao}@gmail.com, zty@powerlaw.ai, {lzy, sms}@tsinghua.edu.cn

## Abstract

We present JEC-QA, the largest question answering dataset in the legal domain, collected from the National Judicial Examination of China. The examination is a comprehensive evaluation of professional skills for legal practitioners. College students are required to pass the examination to be certified as a lawyer or a judge. The dataset is challenging for existing question answering methods, because both retrieving relevant materials and answering questions require the ability of logic reasoning. Due to the high demand of multiple reasoning abilities to answer legal questions, the state-of-the-art models can only achieve about 28% accuracy on JEC-QA, while skilled humans and unskilled humans can reach 81% and 64% accuracy respectively, which indicates a huge gap between humans and machines on this task. We will release JEC-QA and our baselines to help improve the reasoning ability of machine comprehension models. You can access the dataset from <http://jecqa.thunlp.org/>.

## Introduction

Legal Question Answering (LQA) aims to provide explanations, advice or solutions for legal issues. A qualified LQA system can not only provide a professional consulting service for unskilled humans but also help professionals to improve work efficiency and analyze real cases more accurately, which makes LQA an important NLP application in the legal domain. Recently, many researchers attempt to build LQA systems with machine learning techniques (Fawei et al. 2018) and neural network (Do et al. 2017). Despite these efforts in employing advanced NLP models, LQA is still confronted with the following two major challenges. The first is that there is less qualified LQA dataset which limits the research. The second is that the cases and questions in the legal domain are very complex and rigorous. As shown in Table 1, most questions in LQA can be divided into two typical types: the knowledge-driven questions (KD-questions) and case-analysis questions (CA-questions). KD-questions focus on the understanding of spe-

<sup>\*</sup>Indicates equal contribution. The order is determined by dice rolling.

<sup>†</sup>Corresponding author.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

---

**Knowledge-Driven Question:** Which of the following belong to the “property” of Civil Law?

**Option:**

- × A. Trademark.      × B. The star on the sky.
- × C. Gold teeth.      ✓ D. Fish in the pond.

---

**Case-Analysis Question:** Alice owed Bob 3,000 yuan. Alice proposed to pay back with 10,000 yuan of counterfeit money. Bob agreed and accepted it. Which crimes did Alice commit?

**Option:**

- ✓ A. Crime of selling counterfeit money.
  - × B. Crime of using counterfeit money.
  - × C. Crime of embezzlement.
  - × D. Alice did not constitute a crime.
- 

Table 1: Two typical examples of KD-questions and CA-questions in LQA. All examples we show in the paper are translated from Chinese for illustration.

cific legal concepts, while CA-questions concentrate more on the analysis of real cases. Both types of questions require sophisticated reasoning ability and text comprehension ability, which makes LQA a hard task in NLP.

To push forward the development of LQA, we present JEC-QA in this paper, the largest and more challenging LQA dataset. JEC-QA collects questions from the National Judicial Examination of China (NJE) and websites for the examination. NJE is the legal professional certification examination for those who want to be a lawyer or a judge in China. Every year, only around 10% of participants can pass the exam, proving it difficult even for skilled humans.

There are three main properties of JEC-QA: (1) JEC-QA contains 26,365 multiple-choice questions in total, with four options for each question. The number of questions in JEC-QA is 50 times larger than the previous largest LQA dataset (Kim et al. 2016). (2) JEC-QA provides a database including all the legal knowledge required by the examination. The database is collected from the National Unified Legal Professional Qualification Examination Counseling Book and Chinese legal provisions. (3) JEC-QA provides extra labels for questions, including the type of questions (KD-questions or CA-questions) and the reasoning abilities

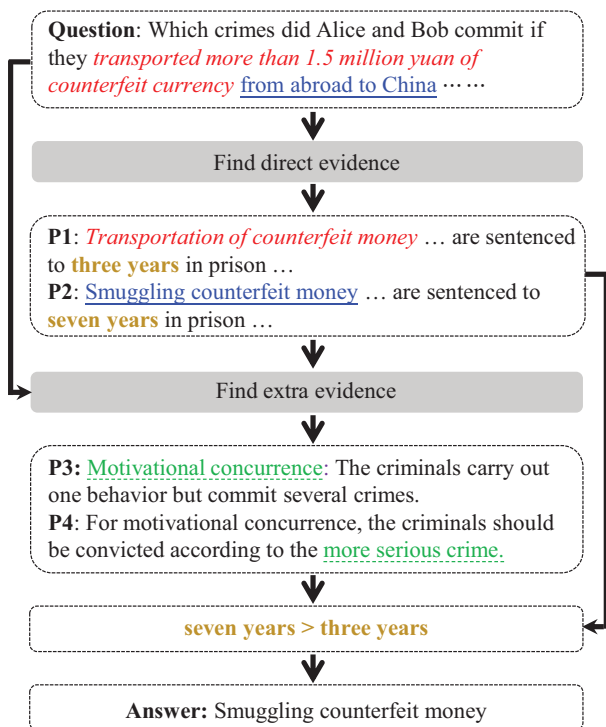


Figure 1: An illustration of the logic that a person answers a question in JEC-QA. **P1** to **P4** are 4 relevant paragraphs retrieved from the legal database. The first two are definitions of two crimes. The last two describe a legal concept and sentencing criterion.

required by the questions. The meta information labeled by skilled humans will be useful for depth-analysis of LQA.

JEC-QA can be addressed following the setting of OpenQA (Chen et al. 2017; Wang et al. 2018b; 2018c; Lin et al. 2018). That is, we need to retrieve relevant articles from the databases and apply reading comprehension models to answer questions. Distinct from existing question answering datasets (Yang, Yih, and Meek 2015; Richardson, Burges, and Renshaw 2013; Hermann et al. 2015; Rajpurkar et al. 2016; Trischler et al. 2016; Lai et al. 2017), JEC-QA requires multiple reasoning abilities to answer the questions including word matching, concept understanding, numerical analysis, multi-paragraph reading, and multi-hop reasoning. The detailed analysis can be found in the section of Reasoning Types.

To get a better understanding of these reasoning abilities, we show a question of JEC-QA in Fig. 1 describing a criminal behavior which results in two crimes. The models must understand “Motivational Concurrence” to reason out extra evidence rather than lexical-level semantic matching. Moreover, the models must have the ability of multi-paragraph reading and multi-hop reasoning to combine the direct evidence and the extra evidence to answer the question, while numerical analysis is also necessary for comparing which crime is more serious. We can see that answering one question will need multiple reasoning abilities in both retrieving

and answering, makes JEC-QA a challenging task.

To investigate the challenges and characteristics of LQA, we design a unified OpenQA framework and implement seven representative neural methods of reading comprehension. By evaluating the performance of these methods on JEC-QA, we show that even the best method can only achieve about 25% and 29% on KD-questions and CA-questions respectively, while skilled humans and unskilled humans can reach 81% and 64% accuracies on JEC-QA. The experimental results show that existing OpenQA methods suffer from the inability of complex reasoning on JEC-QA as they cannot well understand legal concepts and handle multi-hop reasoning.

In summary, JEC-QA is the largest LQA dataset, and it is more challenging compared with existing datasets due to the requirements of multiple reasoning abilities and legal knowledge. JEC-QA will benefit the research of question answering and legal analysis. We also show the performance of existing methods, conduct an in-depth analysis of JEC-QA and outlook the future research direction. You can access the dataset from <http://jecqa.thunlp.org/>.

## Related Work

### Reading Comprehension

There have been numerous reading comprehension datasets proposed in recent years, such as CNN/DailyMail (Hermann et al. 2015), MCTest (Richardson, Burges, and Renshaw 2013), SQuAD (Rajpurkar et al. 2016), WikiQA (Yang, Yih, and Meek 2015) and NewsQA (Trischler et al. 2016). Deep reading comprehension models (Seo et al. 2017; Wang et al. 2017; Wang and Jiang 2016; Dhingra et al. 2017; Yih et al. 2015) have achieved promising results on these early datasets. Besides, recent works like TrivialQA (Joshi et al. 2017), MS-MARCO (Nguyen et al. 2016) and DuReader (He et al. 2018b) contain multiple passages for each question, while RACE (Lai et al. 2017), HotpotQA (Yang et al. 2018) and ARC (Clark et al. 2018) datasets require the ability of reasoning. Based on these datasets, researchers (Wang et al. 2018a; 2018b; 2018d; Clark and Gardner 2018) propose to aggregate information from all passages. These datasets take a step towards a more challenging reading comprehension task, but still have a limitation that the answers can be extracted from the passages directly with semantic matching. As a result, existing RC systems are still lack of reasoning ability and language understanding (Jia and Liang 2017).

### Open-domain Question Answering

OpenQA is first proposed by (Green Jr et al. 1961), which aims to answer questions with external knowledge bases, such as collected documents (Voorhees and others 1999), web-pages (Kwok, Etzioni, and Weld 2001; Chen and Van Durme 2017) or structured knowledge bases (Berant et al. 2013; Bordes et al. 2015; Yu et al. 2017).

Most OpenQA models contain two steps: reading material retrieval and answer extraction/selection (Chen et al. 2017; Dhingra et al. 2017; Cui et al. 2017). Without document-level annotations, they retrieve documents with unsuper-

vised information retrieval methods, e.g., TF-IDF or BM25 retriever. However, these models focus on the lexical similarity between articles and questions rather than semantic relevance. Recent approaches (Lin et al. 2018; Wang et al. 2018c; Clark and Gardner 2018) tend to rerank passages retrieved in the first step and filter out noisy contents. Although these methods can surpass human performance in certain situations, they are still lack of reasoning ability (Rajpurkar, Jia, and Liang 2018).

## Legal Intelligence

Owing to the massive quantity of high-quality textual data in the legal domain, employing NLP techniques to solve legal intelligence problems has been more and more popular in recent years, e.g., generating court views to interpret charge results (Ye et al. 2018), retrieving relevant or similar cases (Chen, Liu, and Ho 2013; Raghav, Reddy, and Reddy 2016), predicting charges or identifying applicable articles (Luo et al. 2017; Hu et al. 2018; He et al. 2018a; Zhong et al. 2018; Xiao et al. 2018; Shen et al. 2018).

Meanwhile, answering legal questions has been a long-standing challenge for applications of legal intelligence. Kim et al.; Kim et al. (2016; 2018) held a legal question answering competition, where rule-based systems (Fawei et al. 2018) and neural models (Do et al. 2017) were applied to this task.

In spite of this, we are still far away from applicable LQA systems, due to the poor performance, reasoning ability, and interpretability. We collect JEC-QA from NJEC, which can serve as a good benchmark of the reasoning ability of legal domain question answering models.

## Dataset Construction and Analysis

### Dataset Construction

**Questions.** We collect 2,700 multiple-choice questions from the 2009 to 2017 national judicial and 30,371 practice exercises from websites. After removing duplicated questions, there are 26,365 questions in JEC-QA.

	KD-questions	CA-questions	Total
Single	4,603	8,738	13,341
Multi	5,158	7,866	13,024
All	9,761	16,604	26,365

Table 2: The statistics of question types in JEC-QA.

Each question in JEC-QA contains a question description and four candidate options. There are single-answer and multi-answer questions in JEC-QA. Meanwhile, we can also classify the questions into Knowledge-Driven Questions (KD-questions) and Case-Analysis Questions (CA-questions). KD-questions pay attention to the definition and interpretation of legal concepts, while CA-questions require analysis for the actual scenarios. Answering both types of questions requires reasoning ability. More detailed statistics of question types are summarized in Table 2.

	Questions	Options	Paragraphs
Count	26,365	105,460	79,433
Average Length	47.01	14.52	58.42
Max Length	547	153	2,738
Vocab Size	29,268	29,987	47,808
Total Vocab Size	70,110		

Table 3: The statistics of questions, options, and reading paragraphs in JEC-QA.

**Database.** As mentioned in introduction, all necessary knowledge for the examination is involved in the National Unified Legal Professional Qualification Examination Counseling Book and Chinese legal provisions. The book contains 15 topics and 215 chapters with highly hierarchically formed contents. To guarantee the retrieval quality, we convert this papyrus book into structured electronic edition manually instead of using OCR (Optical Character Recognition) tools. For Chinese legal provisions, we include 3,382 different legal provisions in our database. The details of the database can be found in Table 3.

### Reasoning Types

We summarize 5 different reasoning types required for answering questions in JEC-QA from JEC-QA and previous works (Lai et al. 2017; Clark et al. 2018), and the examples are shown in Table 4.

(1) **Word Matching.** This is the simplest type of reasoning. The models only need to check which options are matched with the relevant paragraphs and the relevant paragraphs can be easily retrieved by simple search strategies as the contexts are highly consistent. Questions that require this type of reasoning are similar to the ones in traditional reading comprehension datasets.

(2) **Concept Understanding.** As our dataset is built on the legal domain, models need to understand legal concepts to answer these questions. As shown in the 2-nd example in Table 4, models need to understand the meanings of “principal offender” to choose the correct answer.

(3) **Numerical Analysis.** This type of reasoning requires models to perform arithmetic operations. As shown in the 3-rd example in Table 4, models must calculate  $12 \times \frac{1}{3} = 4 < 5$  to answer it.

(4) **Multi-Paragraph Reading.** The settings for previous single-paragraph reading tasks guarantee that enough evidence can be found within one paragraph. However, as shown in the 4-th example in Table 4, specific questions in JEC-QA require reading multiple paragraphs to gather enough evidence, which makes JEC-QA more challenging compared to traditional reading comprehension tasks.

(5) **Multi-Hop Reasoning.** Multi-hop reasoning means that we need multiple steps of logical reasoning to get the answers. Multi-hop reasoning is common in our real lives, but it is hard for existing methods to provide an interpretable reasoning process. Here we show an example of multi-hop reasoning in Fig. 1. Answering this question need to make several steps of reasoning, including concept understanding,

Reasoning Type	KD-Q	CA-Q	All	Examples
Word Matching	<b>65.9%</b>	23.9%	40.5%	<b>Question:</b> Which option is a form of state compensation? <b>Option:</b> <i>Monetary awards</i> <b>Paragraph:</b> <i>Monetary awards</i> is a form of state compensation.
Concept Understanding	36.4%	42.8%	40.2%	<b>Question:</b> Who is the <i>principal offender</i> according to Criminal Law? <b>Option:</b> Bob, the leader of a robbery group, who ordered his subordinates to commit robbery on multiple occasions, but was never personally involved. <b>Paragraph:</b> The <i>principal offender</i> is the person in a group of offenders who leads, organizes, and carries out the main part of a criminal act.
Numerical Analysis	4.6%	14.9%	10.8%	<b>Question:</b> In which of the following circumstances should an extraordinary general meeting of shareholders be convened? <b>Option:</b> The registered capital of the company is <i>12 million yuan</i> , and the unrecovered loss is <i>5 million</i> . <b>Paragraph:</b> In the following circumstances, an extraordinary general meeting of shareholders should be convened: (1) When the unrecovered losses amount to <i>one-third of the total paid-up share capital</i> ; ...
Multi-Paragraph Reading	19.7%	29.4%	25.5%	<b>Question:</b> Which statement is true about corporate crimes? <b>Option:</b> Corporates can be the subject of bank fraud. <b>Paragraph 1:</b> Article 200 of Criminal Law: The punishment of fraud offenses committed by corporates. If a corporate commits any crimes specified in <i>articles 192, 194, or 195 of this section</i> , it shall be fined. <b>Paragraph 2:</b> Article 194 of Criminal Law: <i>Bank fraud</i> ...
Multi-Hop Reasoning	8.33%	<b>66.2%</b>	43.2%	Shown in Fig. 1.

Table 4: Percentages and examples of questions in JEC-QA that require different types of reasoning. We only list one correct option in the table. One question may require multiple reasoning abilities so the sum of percentages is over 100%.

numerical analysis, and multi-paragraph reading. From Table 4, we observe that more than 66% CA-questions require multi-hop reasoning ability, which leads great challenges to existing reading comprehension models.

In conclusion, we summarize that all 5 types of reasoning above are essential for answering questions in JEC-QA and models need to handle these reasoning issues to achieve a promising performance in JEC-QA.

## Experiments

In this section, we conduct detailed experiments and analysis to investigate the performance of existing question answering models on JEC-QA. Following the settings of OpenQA, we first retrieve relevant paragraphs and then employ question answering models to give answers.

### Retrieve Strategy

To retrieve relevant materials from the database, we apply ElasticSearch<sup>1</sup> to build a search engine containing the whole database. As the text materials are hierarchically structured, we store the contents into the search engine with meta-information, such as tags, chapter titles, and section titles. Because different options may focus on various aspects even within the same question, we need to retrieve reading paragraphs for each option separately.

To reduce noisy data and narrow the scope during retrieving, we need to identify the topic (e.g., constitution, criminal law) of the questions. There are 15 topics in total, and we employ 3 representative models, including BERT (Devlin et al. 2018), TextCNN (Kim 2014), and DPCNN (Johnson and

Zhang 2017). From 10,008 labeled instances, we randomly select 1,956 instances for testing and the rest for training. The performance of topic classification is shown in Table 5.

Method	Top-1	Top-2	Top-3
TextCNN	77.97	87.14	91.46
DPCNN	75.16	87.40	92.71
BERT	75.31	88.60	<b>93.10</b>

Table 5: Accuracy (%) of topic classification.

From the experimental results, we can see that the top-1 accuracy of topic classification is unsatisfactory, and the increment is little from top-2 to top-3 (only about 5%). In order to reach a balance of performance and speed, we employ BERT as our topic classifier to select the top-2 relevant topics and retrieve  $K$  most relevant reading paragraphs for each topic. Besides, we also retrieve  $K$  extra reading paragraphs from Chinese legal provisions. In total, we retrieve  $3K$  paragraphs for each option. We choose  $K = 6$  for experiments and we will discuss the reason in Comparative Analysis.

To evaluate the performance of our retrieve strategy, we randomly select 377 questions as Annotation Set for manual annotation and annotate each question with 3 labels, including (1) All Hit (AH): all relevant paragraphs are successfully fetched. (2) Partial Miss (PM): some relevant paragraphs are missing. (3) All Miss (AM): no relevant paragraphs exist in the fetched results.

The evaluation results are listed in Table 6. From this table, we observe that around 46% of the questions can be answered correctly based on retrieved materials. The hit rate of

<sup>1</sup><https://www.elastic.co/>

Type	AH	PM	AM
All questions	45.69	35.77	18.54
KD-questions	<u>59.55</u>	28.09	12.36
CA-questions	38.76	<u>39.61</u>	<u>21.63</u>
Word Matching	<u>62.22</u>	26.67	11.11
Concept Understanding	42.54	35.82	21.64
Numerical Analysis	38.89	33.33	<u>27.78</u>
Multi-Paragraph Reading	38.82	<u>48.24</u>	12.94
Multi-Hop Reasoning	38.89	37.50	23.61

Table 6: Evaluation results (%) of the retrieval strategy.

KD-questions is significantly higher than CA-questions as KD-questions are usually related to specific concepts, which leads to easier retrieval. Among different types of reasoning, the performance in word-matching questions achieves the highest hit rate of 62% as the questions are highly consistent with reading paragraphs. The hit rates for other types achieve substantially lower scores due to the demand for sophisticated reasoning ability.

## Experiment Settings

We employ a controlled experimental setting to ensure a fair comparison among various question answering models. Moreover, we use fastText (Joulin et al. 2017) to pretrain word embeddings on a large-scale legal domain dataset (Xiao et al. 2018). For all models, the dimension of word embeddings is  $w = 200$  and the hidden size of model layers is  $d = 256$ .

As the original tasks of our baselines are various, we design a unified OpenQA framework for them. More specifically, the input for the framework is a triplet  $(q, o, r)$  representing the question, options, and reading paragraphs fetched in the retrieving step.  $q$  is a sequence of words  $(q_1, q_2, \dots, q_{|q|})$ .  $o$  is a tuple of  $n = 4$  word sequences expressed as  $((o_{1,1}, o_{1,2}, \dots, o_{1,|o_1|}), \dots, (o_{n,1}, \dots, o_{n,|o_n|}))$ , corresponding to  $n$  options. Suppose there are  $m = 18$  reading paragraphs for each option, then  $r_{i,j}$  denotes the  $j$ -th reading paragraph for the  $i$ -th option, i.e.,  $r_{i,j} = (r_{i,j,1}, r_{i,j,2}, \dots, r_{i,j,|r_{i,j}|})$ , where  $i \in [1, n]$  and  $j \in [1, m]$ .

For the output, we have two different tasks, i.e., answering single-answer questions and all questions. For single-answer questions, the models need to perform the single-label classification and output a score vector  $score^{single} \in \mathbb{R}^n$  for each question, denoting the probability of each option being correct. For all questions, the models need to output a score vector  $score^{all}$  of length  $2^n - 1$  for each question. Experimental results show that it’s slightly better than using a score vector with length  $n$ . These values denote the probability of each possible combination of options.

Note that some models cannot be directly applied to our task, so we slightly modify them in the following steps:

(1) If the original model only takes the questions and the reading paragraphs as input without options, we apply the model on the concatenation of the question and each option, and obtain a score  $s_i$  for the  $i$ -th option. Then the score vec-

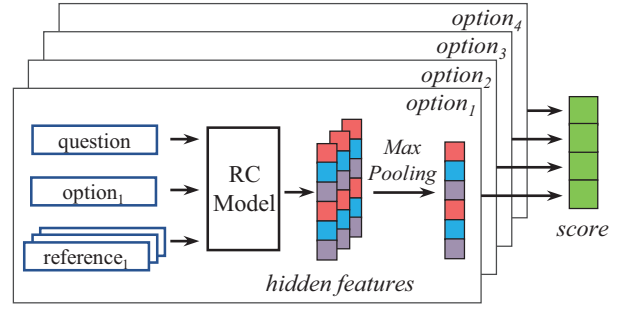


Figure 2: The unified framework for models on JEC-QA.

tor is represented as  $score^{single} = [s_1, s_2, \dots, s_n]$ .

(2) If the original model is designed to extract answers from reading paragraphs, we modify the output layer into a linear layer that outputs the score of the  $i$ -th option,  $s_i$ .

(3) If the original model cannot be applied to multi-paragraph reading task, we apply the model on each reading paragraph of each option separately and the model will output the hidden representation  $h_{i,j} \in \mathbb{R}^d$  for the  $j$ -th reading paragraph of the  $i$ -th option. We then employ max-pooling over all representations from the same option to obtain the hidden representation  $h'_i$  for the  $i$ -th option that we have  $h'_i = [h'_{i,1}, h'_{i,2}, \dots, h'_{i,d}]$  where  $h'_{i,j} = \max(h_{i,k,j} \mid \forall 1 \leq k \leq m)$ . Finally, we pass  $h'_i$  through a linear layer to obtain the score  $s_i$  for the  $i$ -th option.

(4) We add a linear layer with input  $score^{single}$  to obtain  $score^{all}$  for answering all questions.

Besides, we adopt BertAdam (Devlin et al. 2018) for Bert and Adam (Kingma and Ba 2015) for all other models. Meanwhile, for all experiments, we randomly select 20% of the data as the test dataset. You can get more details from the website of the dataset.

## Baselines

We implement 7 representative reading comprehension and question answering models as our baselines, including:

**Co-matching** (Wang et al. 2018a) achieves promising result on the RACE dataset (Lai et al. 2017). The model matches reading paragraphs with questions and options with attention mechanism and uses the attention values to score options. This is a single-paragraph reading comprehension model for single-answer questions.

**BERT** (Devlin et al. 2018) is the model which contains multiple bidirectional Transformer (Vaswani et al. 2017) layers and has been fully pre-trained on large scaled datasets. As a single-paragraph reading comprehension model, **BERT** achieves state-of-the-art performance in most reading comprehension datasets including SQUAD (Rajpurkar et al. 2016). We employ the base form of **BERT** pre-trained on Chinese documents in our experiments.

**SeaReader** (Zhang et al. 2018) is proposed to answer questions in clinical medicine using knowledge extracted from publications in the medical domain. The model extracts information with question-centric attention, document-centric attention, and cross-document attention,

	KD-questions		CA-questions		All	
	Single	All	Single	All	Single	All
Unskilled Humans	76.92	71.11	62.50	58.00	70.00	64.21
Skilled Humans	80.64	77.46	86.84	84.72	84.06	81.12
<b>Co-matching</b> (Wang et al. 2018a)	39.62	<b>25.37</b>	<b>48.91</b>	28.61	<b>46.47</b>	26.06
<b>BERT</b> (Devlin et al. 2018)	38.05	21.13	38.89	23.72	39.56	22.51
<b>SeaReader</b> (Zhang et al. 2018)	39.29	24.11	45.32	26.01	40.50	23.77
<b>Multi-Matching</b> (Tang, Cai, and Zhuo 2019)	<b>41.96</b>	23.63	46.18	<b>29.06</b>	42.98	<b>28.63</b>
<b>CSA</b> (Chen et al. 2019)	32.44	-	34.76	-	21.03	-
<b>CBM</b> (Clark and Gardner 2018)	40.35	22.54	37.37	22.50	38.69	22.53
<b>DSQA</b> (Lin et al. 2018)	34.15	18.41	42.72	23.25	42.63	22.69

Table 7: Evaluation results (accuracy %) of different models on JEC-QA. Results marked “-” indicates that the model cannot converge within 256 epochs.

and then uses a gated layer for denoising.

**Multi-Matching** (Tang, Cai, and Zhuo 2019) employs Evidence-Answer Matching and Question-Passage-Answer Matching module to form matching information, and merges them together to obtain the scores of candidate answers.

**Convolutional Spatial Attention (CSA)** (Chen et al. 2019) first generates enriched representations of passages, candidate answers, and questions with attention mechanism, and then applies CNN-MaxPooling operation to summarize adjacent attention information.

**Confidence-based Model (CBM)** (Clark and Gardner 2018) is a simple and effective method for multi-paragraph reading comprehension task. They propose a pipeline method for single-paragraph reading comprehension and apply a confidence-based method to adapt the model to the multi-paragraph setting.

**Distantly Supervised Question Answering (DSQA)** (Lin et al. 2018) is an effective method for open-domain question answering, which decomposes the QA process into three steps: filter out noisy documents, extract correct answers and select the best answer.

## Experimental Results

We evaluate the performance of all models on JEC-QA, with settings of single-answer question and all question answering. Besides, we also evaluate the performance in KD-questions and CA-questions separately. In addition, we evaluate the performance of skilled and unskilled humans. Humans read the same paragraphs fetched by the search strategy as models do. Unskilled humans are those who do not have legal experience while skilled humans are those in legal professions. The experimental results are shown in Table 7.

From these results, we observe that even the best-performed model can only achieve an accuracy of 28.63% on all questions, while there is still a huge gap to 64% accuracy for unskilled humans. We should note that unskilled humans read the same reading materials as models and they have no advanced knowledge about legal questions, so the gap mainly comes from the insufficiency of model reasoning ability. Meanwhile, compared with skilled humans, unskilled humans perform significantly worse than skilled humans on CA-questions. The reason is that retrieved reading

	KD-Q	CA-Q	All
Word Matching	20.20	28.00	31.91
Concept Understanding	30.35	20.83	28.24
Numerical Analysis	16.67	25.71	30.00
Multi-Paragraph Reading	23.33	19.44	30.51
Multi-Hop Reasoning	25.00	18.62	30.30
All Hit	22.34	24.47	31.71
Partial Miss	29.73	24.00	26.76
All Miss	21.05	16.36	29.79

Table 8: Performance of Co-matching on different questions.

paragraphs are insufficient to provide enough evidence, as shown in Table 6, so the gap between unskilled and skilled humans mainly comes from the quality of retrieval.

Comparing the performance between KD-questions and CA-questions, we reveal that most models achieve better performance on CA-questions. Although a higher proportion of CA-questions require multi-hop reasoning ability, the concepts in CA-questions are always simpler ones, e.g., robbery, theft, or murder. The results also demonstrate that existing methods performs poorly in concept comprehension.

## Comparative Analysis

We also perform a deeper analysis on the well-performed **Co-matching** by evaluating it on Annotation Set and the experimental results are listed in Table 8. From the results we can see that existing methods can only answer about 32% of questions correctly even when there is enough evidence in reading paragraphs, which means that the models cannot understand the reading materials at all. Moreover, we can see that the model performs extremely bad on multi-paragraph reading and multi-hop reasoning questions of CA-questions. It means existing models cannot do multi-paragraph reading and multi-hop reasoning on real cases properly.

Besides, we also perform experiments with different value of  $K$  on single KD-questions, and the experimental results are shown in Table 9. From the results, we can see that more reading paragraphs cannot help the models to answer the questions better, as important articles have already been

etched even  $K$  is small. It proves that the bad performance of models is because the insufficiency of reasoning ability rather than the quality of retrieval. As a larger value of  $K$  cannot help with the accuracy, we select  $K = 6$  to reach a balance of speed and performance.

$K =$	1	3	6	12	18	24
Accuracy	30.1	37.9	39.6	40.7	40.7	40.7

Table 9: Performance of Co-matching with different  $K$ .

## Case Study

As shown in Table 10, we select an example to give an intuitive illustration on dealing with multi-hop reasoning. For most reading comprehension models, they choose all the options as their answers. Even without reading the statement, we can find that the option D conflicts with the other three options. Existing methods cannot handle conflicting options. Moreover, if we ignore option D, these models still choose all the remaining options, while the correct answer only contains option C. The models can easily find the evidence of option A, B, C from the statement with one-hop reasoning. However, if we read the related paragraphs, we will find the fact that Bob is under the age of 16, which will filter out the options A and B. We can learn that existing reading comprehension models already have the ability of one-hop reasoning, but multi-hop reasoning is still challenging for them.

**Question:** Bob is a male **born on February 27, 1987**. Bob stole from Alice a total of 5,000 yuan in cash, one laptop (worth 13,000 yuan), and other small jewelry on **February 27, 2003**. While Bob was climbing back over the wall, Bob was seen by Catherine. To escape, Bob quickly took a dagger from his pocket and stabbed in Catherine’s heart, *killing Catherine*. So how should Bob’s behavior be handled?

**Options:**

- × (A). The crime of robbery.
- × (B). The crime of theft.
- ✓ (C). The ground of intentional homicide.
- × (D). Bob does not constitute a crime.

**Paragraphs:**

1. A person who has **reached the age of 14 and under 16** will not constitute the crime of robbery and theft.
2. Calculation of age. ... For example, you are **14 years old from the next day after your 14-th birthday**.

Table 10: A multi-hop reasoning example.

## Conclusion

In this work, we present JEC-QA as a new and challenging dataset for LQA, and JEC-QA is the largest dataset in LQA. Both retrieving documents and answering questions in JEC-QA require multiple types of reasoning ability, and our experimental results show that existing state-of-the-art models cannot perform well on JEC-QA. We hope our JEC-QA can benefit researchers on improving the reasoning ability

of reading comprehension and QA models, and also making advances for legal question answering.

In the future, we will explore how to improve the reasoning ability of question answering model and integrate legal knowledge into question answering, which are necessary for answering questions in JEC-QA.

## Acknowledgements

This work is supported by the National Key Research and Development Program of China (No. 2018YFC0831900) and the National Natural Science Foundation of China (NSFC No. 61572273, 61661146007).

## References

- Berant, J.; Chou, A.; Frostig, R.; and Liang, P. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of EMNLP*.
- Bordes, A.; Usunier, N.; Chopra, S.; and Weston, J. 2015. Large-scale simple question answering with memory networks. *arXiv preprint arXiv:1506.02075*.
- Chen, T., and Van Durme, B. 2017. Discriminative information retrieval for question answering sentence selection. In *Proceedings of EACL*.
- Chen, D.; Fisch, A.; Weston, J.; and Bordes, A. 2017. Reading wikipedia to answer open-domain questions. In *Proceedings of ACL*.
- Chen, Z.; Cui, Y.; Ma, W.; Wang, S.; and Hu, G. 2019. Convolutional spatial attention model for reading comprehension with multiple-choice questions. In *Proceedings of AAAI*.
- Chen, Y.-L.; Liu, Y.-H.; and Ho, W.-L. 2013. A text mining approach to assist the general public in the retrieval of legal documents. *Journal of ASIS&T* 64(2):280–290.
- Clark, C., and Gardner, M. 2018. Simple and effective multi-paragraph reading comprehension. In *Proceedings of ACL*.
- Clark, P.; Cowhey, I.; Etzioni, O.; Khot, T.; Sabharwal, A.; Schoenick, C.; and Tafjord, O. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Cui, Y.; Chen, Z.; Wei, S.; Wang, S.; Liu, T.; and Hu, G. 2017. Attention-over-attention neural networks for reading comprehension. In *Proceedings of ACL*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dhingra, B.; Liu, H.; Yang, Z.; Cohen, W. W.; and Salakhutdinov, R. 2017. Gated-attention readers for text comprehension. In *Proceedings of ACL*.
- Do, P.-K.; Nguyen, H.-T.; Tran, C.-X.; Nguyen, M.-T.; and Nguyen, M.-L. 2017. Legal question answering using ranking svm and deep convolutional neural network. *arXiv preprint arXiv:1703.05320*.
- Fawei, B.; Pan, J. Z.; Kollingbaum, M.; and Wyner, A. Z. 2018. A methodology for a criminal law and procedure ontology for legal question answering. In *Proceedings of JIST*.
- Green Jr, B. F.; Wolf, A. K.; Chomsky, C.; and Laughery, K. 1961. Baseball: an automatic question-answerer. In *Proceedings of IRE-AIEE-ACM*.
- He, C.; Peng, L.; Le, Y.; and He, J. 2018a. Secaps: A sequence enhanced capsule model for charge prediction. *arXiv preprint arXiv:1810.04465*.

- He, W.; Liu, K.; Liu, J.; Lyu, Y.; Zhao, S.; Xiao, X.; Liu, Y.; Wang, Y.; Wu, H.; She, Q.; et al. 2018b. Dureader: a chinese machine reading comprehension dataset from real-world applications. In *Proceedings of ACL workshop*.
- Hermann, K. M.; Kocisky, T.; Grefenstette, E.; Espeholt, L.; Kay, W.; Suleyman, M.; and Blunsom, P. 2015. Teaching machines to read and comprehend. In *Proceedings of NIPS*.
- Hu, Z.; Li, X.; Tu, C.; Liu, Z.; and Sun, M. 2018. Few-shot charge prediction with discriminative legal attributes. In *Proceedings of COLING*.
- Jia, R., and Liang, P. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of EMNLP*.
- Johnson, R., and Zhang, T. 2017. Deep pyramid convolutional neural networks for text categorization. In *Proceedings of ACL*.
- Joshi, M.; Choi, E.; Weld, D. S.; and Zettlemoyer, L. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*.
- Joulin, A.; Grave, E.; Bojanowski, P.; Douze, M.; Jégou, H.; and Mikolov, T. 2017. Fasttext. zip: Compressing text classification models. In *Proceedings of ICLR*.
- Kim, M.-Y.; Goebel, R.; Kano, Y.; and Satoh, K. 2016. Coliee-2016: evaluation of the competition on legal information extraction and entailment. In *Proceedings of JURISIN*.
- Kim, M.-Y.; Lu, Y.; Rabelo, J.; and Goebel, R. 2018. Coliee-2018: Evaluation of the competition on case law information extraction and entailment.
- Kim, Y. 2014. Convolutional neural networks for sentence classification. In *Proceedings of EMNLP*.
- Kingma, D., and Ba, J. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR*.
- Kwok, C.; Etzioni, O.; and Weld, D. S. 2001. Scaling question answering to the web. *ACM Transactions on Information Systems*.
- Lai, G.; Xie, Q.; Liu, H.; Yang, Y.; and Hovy, E. 2017. RACE: Large-scale reading comprehension dataset from examinations. In *Proceedings of EMNLP*.
- Lin, Y.; Ji, H.; Liu, Z.; and Sun, M. 2018. Denoising distantly supervised open-domain question answering. In *Proceedings of ACL*.
- Luo, B.; Feng, Y.; Xu, J.; Zhang, X.; and Zhao, D. 2017. Learning to predict charges for criminal cases with legal basis. In *Proceedings of EMNLP*.
- 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*.
- Raghav, K.; Reddy, P. K.; and Reddy, V. B. 2016. Analyzing the extraction of relevant legal judgments using paragraph-level and citation information. In *Proceedings of ECAI*.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of EMNLP*.
- Rajpurkar, P.; Jia, R.; and Liang, P. 2018. Know what you don't know: Unanswerable questions for squad. In *Proceedings of ACL*.
- Richardson, M.; Burges, C. J.; and Renshaw, E. 2013. Mctest: A challenge dataset for the open-domain machine comprehension of text. In *Proceedings of EMNLP*.
- Seo, M.; Kembhavi, A.; Farhadi, A.; and Hajishirzi, H. 2017. Bidirectional attention flow for machine comprehension. In *Proceedings of ICLR*.
- Shen, Y.; Sun, J.; Li, X.; Zhang, L.; Li, Y.; and Shen, X. 2018. Legal article-aware end-to-end memory network for charge prediction. In *Proceedings of ICCSE*.
- Tang, M.; Cai, J.; and Zhuo, H. H. 2019. Multi-matching network for multiple choice reading comprehension. In *Proceedings of AAAI*.
- Trischler, A.; Wang, T.; Yuan, X.; Harris, J.; Sordani, A.; Bachman, P.; and Suleman, K. 2016. Newsqa: A machine comprehension dataset. *arXiv preprint arXiv:1611.09830*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Proceedings of NIPS*.
- Voorhees, E. M., et al. 1999. The trec-8 question answering track report. In *Proceedings of Trec*.
- Wang, S., and Jiang, J. 2016. Machine comprehension using match-lstm and answer pointer. *arXiv preprint arXiv:1608.07905*.
- 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 ACL*.
- Wang, S.; Yu, M.; Chang, S.; and Jiang, J. 2018a. A co-matching model for multi-choice reading comprehension. In *Proceedings of ACL*.
- Wang, S.; Yu, M.; Guo, X.; Wang, Z.; Klinger, T.; Zhang, W.; Chang, S.; Tesauro, G.; Zhou, B.; and Jiang, J. 2018b. R<sup>3</sup>: Reinforced reader-ranker for open-domain question answering. In *Proceedings of AAAI*.
- Wang, S.; Yu, M.; Jiang, J.; Zhang, W.; Guo, X.; Chang, S.; Wang, Z.; Klinger, T.; Tesauro, G.; and Campbell, M. 2018c. Evidence aggregation for answer re-ranking in open-domain question answering. In *Proceedings of ICLR*.
- Wang, Y.; Liu, K.; Liu, J.; He, W.; Lyu, Y.; Wu, H.; Li, S.; and Wang, H. 2018d. Multi-passage machine reading comprehension with cross-passage answer verification. In *Proceedings of ACL*.
- Xiao, C.; Zhong, H.; Guo, Z.; Tu, C.; Liu, Z.; Sun, M.; Feng, Y.; Han, X.; Hu, Z.; Wang, H.; et al. 2018. Cail2018: A large-scale legal dataset for judgment prediction. *arXiv preprint arXiv:1807.02478*.
- Yang, Z.; Qi, P.; Zhang, S.; Bengio, Y.; Cohen, W.; Salakhutdinov, R.; and Manning, C. D. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of EMNLP*.
- Yang, Y.; Yih, W.-t.; and Meek, C. 2015. Wikiqa: A challenge dataset for open-domain question answering. In *Proceedings of EMNLP*.
- Ye, H.; Jiang, X.; Luo, Z.; and Chao, W. 2018. Interpretable charge predictions for criminal cases: Learning to generate court views from fact descriptions. In *Proceedings of NAACL*.
- Yih, W.-t.; Chang, M.-W.; He, X.; and Gao, J. 2015. Semantic parsing via staged query graph generation: Question answering with knowledge base. In *Proceedings of ACL*.
- Yu, M.; Yin, W.; Hasan, K. S.; Santos, C. d.; Xiang, B.; and Zhou, B. 2017. Improved neural relation detection for knowledge base question answering. In *Proceedings of ACL*.
- Zhang, X.; Wu, J.; He, Z.; Liu, X.; and Su, Y. 2018. Medical exam question answering with large-scale reading comprehension. In *Proceedings of AAAI*.
- Zhong, H.; Zhipeng, G.; Tu, C.; Xiao, C.; Liu, Z.; and Sun, M. 2018. Legal judgment prediction via topological learning. In *Proceedings of EMNLP*.