

# When Retriever-Reader Meets Scenario-Based Multiple-Choice Questions

Zixian Huang and Ao Wu and Yulin Shen and Gong Cheng and Yuzhong Qu

State Key Laboratory for Novel Software Technology, Nanjing University, China

{zixianhuang, awu, ylshen}@smail.nju.edu.cn

{gcheng, yzqu}@nju.edu.cn

## Abstract

Scenario-based question answering (SQA) requires retrieving and reading paragraphs from a large corpus to answer a question which is contextualized by a long scenario description. Since a scenario contains both keyphrases for retrieval and much noise, retrieval for SQA is extremely difficult. Moreover, it can hardly be supervised due to the lack of relevance labels of paragraphs for SQA. To meet the challenge, in this paper we propose a joint retriever-reader model called JEEVES where the retriever is implicitly supervised only using QA labels via a novel word weighting mechanism. JEEVES significantly outperforms a variety of strong baselines on multiple-choice questions in three SQA datasets.

## 1 Introduction

Scenario-based question answering (SQA) is the task of answering a question which is contextualized by a *scenario* describing a concrete case (Lally et al., 2017), e.g., predicting a judgment based on a legal scenario (Xu et al., 2020), answering a multiple-choice question in China’s national college entrance examination (i.e., GaoKao) based on a geographical scenario (Li et al., 2021). SQA is a challenging task as it combines the difficulties of open-domain question answering (ODQA) (Chen et al., 2017; Kwiatkowski et al., 2019) and machine reading comprehension (MRC) (Rajpurkar et al., 2016; Lai et al., 2017). Indeed, SQA not only relies on accurately *retrieving* relevant paragraphs from a large corpus, but also requires thoroughly *reading* the retrieved paragraphs and the provided scenario to fuse their knowledge and infer an answer.

For example, Figure 1 shows a scenario and a multiple-choice question sampled from geography exams in GaoKao. For machines to answer such a question, it would be insufficient to only read the given scenario but is necessary to retrieve from a corpus (e.g., textbooks, Wikipedia) to acquire supporting knowledge. In this example, the scenario

**Scenario:** Tian Xiao in his essay *My Hometown* wrote:  
*In an April morning, the vast expanse of farmland looked endless. Boats sailed through dense networks of rivers. A slight breeze blew. Green rice seedlings in the farmland danced against the breeze. Alongside the farmland, girls harvesting mulberries sang and laughed. Fishes in the pond frequently leaped out of water. Alongside the pond, sugar cane seedlings thrived. Flowers in fruit trees smelt nice. Spring was much in the air.*

**Question:** The climate type in the described area is

**Options:** (A) temperate monsoon climate  
(B) temperate continental climate  
(C) tropical rainforest climate  
(D) subtropical monsoon climate (Correct Answer)

**Pond Aquaculture on Sugarcane Farm**  
Pond aquaculture on sugarcane farm refers to an aquaculture production facility which also grows sugar cane alongside fish ponds... Pond mud is used as fertilizer, while fishes are fed with sugarcane leaves and tops.

**Pearl River Delta**  
... It is mostly located in the tropics and is with a subtropical monsoon climate... It has a very high river network density... Farmers in the area have developed... pond aquaculture on sugarcane farms...

Figure 1: **Top:** a scenario and a multiple-choice question about geography; in the scenario we highlight some keyphrases for retrieval to support answering the question. **Bottom:** some supporting paragraphs in two articles retrieved from an online encyclopedia. **Arrows:** two reasoning paths that connect the area described in the scenario with a similar area in the real world.

gives clues about the location of the described area. The keyphrases we highlight in the scenario connect the described area with the Pearl River Delta via two reasoning paths. The delta has a subtropical monsoon climate, indicating the correct answer. This multi-hop reasoning process fuses the scenario and the retrieved supporting paragraphs.

**Research Challenge.** SQA, similar to ODQA, is commonly solved by a retriever-reader framework. However, retrieval for SQA is fundamentally more difficult than for ODQA because a scenario typically gives a long paragraph containing both *keyphrases for retrieval* and much *noise*, as illustrated in Figure 1. Hence, in SQA, the paragraphs found by unsupervised retrievers such as BM25 are often not truly useful for answering the question. Retrieval errors have become a bottleneck in SQA (Cheng et al., 2016; Huang et al., 2019).

Supervised retrievers may be more powerful but they cannot directly apply to SQA due to the *lack of relevance labels of paragraphs* for training. Only QA labels (i.e., correct answers) are available. Re-

cent dense passage retrievers (DPR) (Karpukhin et al., 2020) can be distantly supervised by QA labels, using paragraphs that contain the gold-standard answer as positives. While such distant supervision is effective for ODQA where answers are short, it could not generate high-quality positives for most questions in SQA where the gold-standard answer is sentence-long and *not explicitly contained in any paragraph*. As a result, DPR shows unsatisfying performance on SQA datasets, as our experimental results will show. How can we effectively supervise a retriever for SQA?

**Our Approach—JEEVES.** To meet the challenge, we propose a novel retriever-reader method called JEEVES where we employ only QA labels to *implicitly* train our retriever. It is realized by learning a *word weighting* for distinguishing keyphrases from noise in the scenario to improve the accuracy of retrieval and hence of the downstream QA which is supervised. Since our retriever and reader are both supervised by QA labels, we train them jointly to further improve the overall accuracy.

We summarize our contributions as follows.

- We propose to implicitly supervise a retriever for answering multiple-choice questions in SQA by learning a word weighting.
- We present JEEVES, a new end-to-end model for SQA where the retriever and the reader are jointly trained only with QA labels.
- JEEVES significantly outperforms state-of-the-art methods on three SQA datasets, including a new dataset we construct.

Our code and data are on GitHub: <https://github.com/nju-websoft/Jeeves-GKMC>.

**Outline.** We discuss related work in Section 2, describe JEEVES in Section 3, present experiments in Section 4, and conclude the paper in Section 5.

## 2 Related Work

**SQA** has found application in many domains (Lally et al., 2017; Zhong et al., 2018; Chalkidis et al., 2019; Yang et al., 2019; Xu et al., 2020; Cheng et al., 2016; Huang et al., 2019; Li et al., 2021). Among others, Huang et al. (2019) published the GeoSQA dataset containing multiple-choice questions in high-school geography exams, and they evaluated a variety of methods, mostly

adopting a retriever-reader framework. All the tested methods performed close to random guess. The authors pointed out that a major cause of error is imprecise retrieval due to the long description of a scenario which is critical to retrieval but noisy.

JEEVES mainly aims at addressing this issue and it outperformed state-of-the-art methods on three SQA datasets in our experiments.

**Retrievers** for SQA as well as ODQA are mainly unsupervised (e.g., BM25) due to the lack of relevance labels of paragraphs. Although word matching based retrievers are still useful (Lee et al., 2019), they could not effectively exploit long and noisy scenario descriptions in SQA and became a bottleneck (Cheng et al., 2016; Huang et al., 2019). Supervised retrievers such as rerankers (Matsubara et al., 2020) and dense retrievers (Lee et al., 2019; Karpukhin et al., 2020; Khattab and Zaharia, 2020) need labeled paragraphs for training. While dense retrievers have been effectively trained with labels from distant supervision on ODQA datasets, in our experiments they were less effective on SQA datasets such as GeoSQA where answers (i.e., options of multiple-choice questions) are sentence-long and often not explicitly contained in any paragraph. As a result, distant supervision by exact matching failed to generate positives, while approximate matching generated false positives.

Another way to handle the lack of labels is by transfer learning. For example, AR (Pirtoaca et al., 2019) trains its retriever with labeled paragraphs from another similar dataset. However, we are not aware of any datasets that contain labeled paragraphs and are sufficiently similar to SQA datasets like GeoSQA. Training with labeled paragraphs from a very different dataset would influence the accuracy of the learned retriever on the target dataset. Indeed, AR showed unsatisfying performance on SQA datasets in our experiments.

JEEVES does not rely on distant supervision but only employs QA labels in the target SQA dataset to implicitly yet effectively train a retriever.

**Readers** have been extensively studied in MRC research. State-of-the-art methods use pre-trained language models (PLMs) (Jin et al., 2020; Zhang et al., 2020a,b; Liu et al., 2020b). Some reader-retriever frameworks for ODQA jointly train their retriever and reader (Wang et al., 2018; Lee et al., 2019). JEEVES incorporates a reader which also builds on PLMs and fuses the knowledge in multi-

ple paragraphs. JEEVES also performs joint training. These design choices showed effectiveness in our ablation study, although they are not included in our main research contributions.

### 3 Approach

#### 3.1 Task Definition

We focus on scenario-based multiple-choice questions. Given a scenario description  $S$ , a question  $Q$  has a set of  $m$  options  $\mathbf{O} = \{O_1, \dots, O_m\}$  where exactly one option is the correct answer. For each option  $O_i \in \mathbf{O}$ , we refer to the concatenation of  $S$ ,  $Q$ , and  $O_i$  as an enriched option  $\hat{O}_i$ . Observe that to answer  $Q$ , we commonly need to fuse the description in  $S$  with knowledge in some supporting paragraphs retrieved from a corpus  $\mathbf{P}$ .

#### 3.2 Overview of JEEVES

Figure 2 presents an overview of JEEVES, which adopts a retriever-reader framework.

**Retriever.** For each enriched option  $\hat{O}_i$ , we consider all the potentially relevant paragraphs  $\mathbf{P}_i^{\text{raw}}$  in the corpus. Since  $\mathbf{P}_i^{\text{raw}}$  may have a very large size, for the paragraphs in  $\mathbf{P}_i^{\text{raw}}$  we only generate their sparse representations  $\mathbf{B}_i$  over the words in  $\hat{O}_i$ , but we weight these words based on the dense representation  $\mathbf{H}_i$  of  $\hat{O}_i$ . We combine sparse representations  $\mathbf{B}_i$  and word weights  $\mathbf{w}_i$  to score each paragraph in  $\mathbf{P}_i^{\text{raw}}$  and retrieve  $k$  top-ranked paragraphs  $\mathbf{P}_i^{\text{top}}$  to be fed into the reader. In particular, we leverage the computed paragraph scores  $\mathbf{z}_i^{\text{spa}}$  to compute a score  $s_i^{\text{spa}}$  for option  $O_i$ , which in the training phase is compared with QA labels to implicitly supervise the retriever.

**Reader.** Since  $\mathbf{P}_i^{\text{top}}$  has a small size of  $k$ , for the paragraphs in  $\mathbf{P}_i^{\text{top}}$  we can afford to generate their dense representations  $\mathbf{G}_i$ . We use  $\mathbf{G}_i$  to re-score each paragraph in  $\mathbf{P}_i^{\text{top}}$ , and we choose  $k'$  top-ranked paragraphs  $\mathbf{P}_i^{\text{fus}}$  to fuse their dense representations into  $\mathbf{f}_i$  and compute a score  $s_i^{\text{fus}}$  for option  $O_i$ . We also leverage the re-computed paragraph scores  $\mathbf{z}_i^{\text{den}}$  to compute a score  $s_i^{\text{den}}$  for  $O_i$ . Finally,  $s_i^{\text{spa}}$ ,  $s_i^{\text{den}}$ , and  $s_i^{\text{fus}}$  are combined to score  $O_i$ .

#### 3.3 Retriever

Our retriever for SQA is implicitly supervised by QA labels via a novel word weighting mechanism.

**Paragraph Representation (BoW Vectors).** For each enriched option  $\hat{O}_i$ , from the corpus  $\mathbf{P}$  we

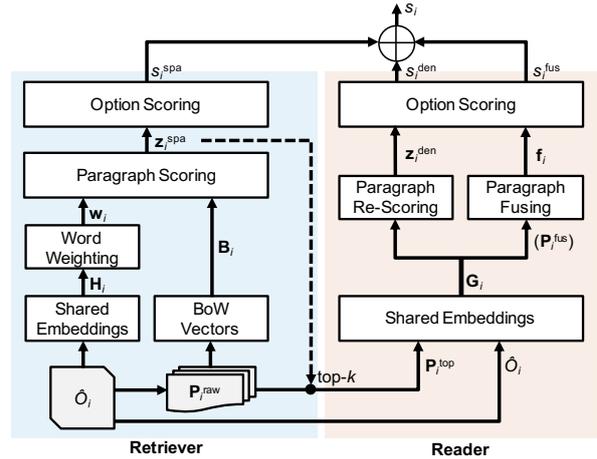


Figure 2: Overview of JEEVES.

take all the potentially relevant paragraphs  $\mathbf{P}_i^{\text{raw}}$ : those containing any non-stopword in  $\hat{O}_i$ . Since  $\mathbf{P}_i^{\text{raw}}$  may have a large size, for each paragraph  $p_l \in \mathbf{P}_i^{\text{raw}}$  we inexpensively generate its sparse representation  $\mathbf{b}_l$  capturing its lexical features. Specifically, we adopt a bag-of-words (BoW) model. Let  $n_i$  be the number of unique words in  $\hat{O}_i$ . We define  $\mathbf{b}_l$  as an  $n_i$ -dimensional vector where each value is the BM25 score of  $p_l$  given the corresponding word in  $\hat{O}_i$  as a query. We pack together such vectors for all the paragraphs in  $\mathbf{P}_i^{\text{raw}}$  into a matrix  $\mathbf{B}_i$ :

$$\mathbf{B}_i = [\mathbf{b}_1; \mathbf{b}_2; \dots]. \quad (1)$$

#### Option Representation (Shared Embeddings).

For  $\hat{O}_i$  we generate its dense representation capturing its semantic features by feeding it into a PLM (Cui et al., 2019; Sun et al., 2019) which will be shared with our reader:

$$\text{PLM}([\text{CLS}] S [\text{SEP}] Q [\text{SEP}] O_i [\text{SEP}]). \quad (2)$$

Recall that  $\hat{O}_i$  contains  $n_i$  unique words. If a word appears in multiple positions in  $\hat{O}_i$ , we will aggregate their dense representations in Eq. (2) by max pooling. We pack together the aggregate vectors for all the  $n_i$  unique words in  $\hat{O}_i$  into a matrix  $\mathbf{H}_i$ :

$$\mathbf{H}_i = [\mathbf{h}_1; \dots; \mathbf{h}_{n_i}]. \quad (3)$$

**Word Weighting.** Some words in  $\hat{O}_i$  are keyphrases for retrieval while others are noise. We learn to weight each word representing its salience. Specifically, for all the  $n_i$  unique words in  $\hat{O}_i$ , we feed their dense representations  $\mathbf{H}_i$  into two dense layers (called *word weighting network*) to output

$n_i$  values normalized by softmax as word weights:

$$\begin{aligned} \mathbf{w}_i &= \text{softmax}([w_{i,1}; \dots; w_{i,n_i}]^\top), \text{ where} \\ w_{i,j} &= \text{linear}(\tanh(\text{linear}(\mathbf{h}_j))) \text{ for } 1 \leq j \leq n_i. \end{aligned} \quad (4)$$

**Paragraph Scoring.** For all the paragraphs in  $\mathbf{P}_i^{\text{raw}}$ , we compute weighted sums over their sparse representations  $\mathbf{B}_i$  to obtain their scores  $\mathbf{z}_i^{\text{spa}}$ :

$$\mathbf{z}_i^{\text{spa}} = \mathbf{B}_i^\top \times \mathbf{w}_i. \quad (5)$$

We retrieve  $k$  top-ranked paragraphs  $\mathbf{P}_i^{\text{top}} \subseteq \mathbf{P}_i^{\text{raw}}$  according to  $\mathbf{z}_i^{\text{spa}}$  and feed  $\mathbf{P}_i^{\text{top}}$  into the reader.

**Option Scoring.** An enrich option stated explicitly in a paragraph is likely the correct answer. Paragraph scores  $\mathbf{z}_i^{\text{spa}}$  represent the confidence that such a statement is found (Clark et al., 2016). Therefore, we leverage  $\mathbf{z}_i^{\text{spa}}$  to compute a score  $s_i^{\text{spa}}$  for option  $O_i$  by feeding  $\mathbf{z}_i^{\text{spa}}$  into two dense layers:

$$s_i^{\text{spa}} = \text{linear}(\tanh(\text{linear}(\mathbf{z}_i^{\text{spa}}))), \quad (6)$$

where we trim  $\mathbf{z}_i^{\text{spa}}$  to its  $\tau$  largest values. In this way, retrieval results  $\mathbf{P}_i^{\text{top}}$  are connected with QA results  $s_i^{\text{spa}}$  to allow implicit supervision.

**Loss Function.** Given a labeled correct answer  $O_j \in \mathbf{O}$ , we calculate the cross-entropy loss:

$$\mathcal{L}^{\text{RETR}} = -\log \frac{\exp(s_j^{\text{spa}})}{\sum_{O_i \in \mathbf{O}} \exp(s_i^{\text{spa}})}. \quad (7)$$

Indeed, we employ only QA labels in Eqs. (6)(7) to implicitly supervise the retriever in Eq. (5) by training the word weighting network in Eq. (4).

### 3.4 Reader

Our reader for SQA uses both separate and fused dense representations of paragraphs.

**Paragraph Representation (Shared Embeddings).** Since  $\mathbf{P}_i^{\text{top}}$  contains a small number of  $k$  paragraphs, for each paragraph  $p_l \in \mathbf{P}_i^{\text{top}}$  we generate its dense representation capturing its semantic features by reusing the shared PLM in Eq. (2):

$$\begin{aligned} [\mathbf{g}_{i,l,1}; \mathbf{g}_{i,l,2}; \dots] &= \\ \text{PLM}([\text{CLS}] p_l [\text{SEP}] S [\text{SEP}] Q [\text{SEP}] O_i [\text{SEP}]). \end{aligned} \quad (8)$$

We take  $\mathbf{g}_{i,l,1}$ , the dense representation of the [CLS] token, as an aggregate representation of the entire sequence. We pack together such vectors for all the  $k$  paragraphs in  $\mathbf{P}_i^{\text{top}}$  into a matrix  $\mathbf{G}_i$ :

$$\mathbf{G}_i = [\mathbf{g}_{i,1,1}; \dots; \mathbf{g}_{i,k,1}]. \quad (9)$$

**Paragraph Re-Scoring.** For all the  $k$  paragraphs in  $\mathbf{P}_i^{\text{top}}$ , we feed their dense representations  $\mathbf{G}_i$  into a dense layer to re-compute their scores  $\mathbf{z}_i^{\text{den}}$ :

$$\mathbf{z}_i^{\text{den}} = [z_{i,1}^{\text{den}}; \dots; z_{i,k}^{\text{den}}]^\top, \text{ where } z_{i,l}^{\text{den}} = \text{linear}(\mathbf{g}_{i,l,1}). \quad (10)$$

These scores are expected to be more accurate than those computed by Eq. (5) based on sparse representations. Therefore, let  $\mathbf{P}_i^{\text{fus}} \subseteq \mathbf{P}_i^{\text{top}}$  be  $k'$  top-ranked paragraphs according to  $\mathbf{z}_i^{\text{den}}$ . Below we will only fuse this subset of paragraphs for both effectiveness and efficiency considerations.

**Paragraph Fusion.** Fusing the knowledge in multiple paragraphs is needed for SQA. We fuse the dense representations of all the  $k'$  paragraphs in  $\mathbf{P}_i^{\text{fus}}$  to enhance their synergy. Specifically, we firstly perform *intra-paragraph fusion*. For each paragraph  $p_l \in \mathbf{P}_i^{\text{fus}}$ , we pack together the vectors in Eq. (8) into four matrices:  $\mathbf{G}_{i,l}^{\text{P}}$  for  $p_l$ ,  $\mathbf{G}_{i,l}^{\text{S}}$  for  $S$ ,  $\mathbf{G}_{i,l}^{\text{Q}}$  for  $Q$ , and  $\mathbf{G}_{i,l}^{\text{O}}$  for  $O_i$ . For each pair of matrices such as  $(\mathbf{G}_{i,l}^{\text{P}}, \mathbf{G}_{i,l}^{\text{S}})$ , we fuse them into a vector  $\mathbf{f}_{i,l}^{\text{PS}}$  by dual-attention (Liu et al., 2020a):

$$\begin{aligned} (\hat{\mathbf{G}}_{i,l}^{\text{P}}, \hat{\mathbf{G}}_{i,l}^{\text{S}}) &= \text{DualAttention}(\mathbf{G}_{i,l}^{\text{P}}, \mathbf{G}_{i,l}^{\text{S}}), \\ \mathbf{f}_{i,l}^{\text{P}} &= \text{max-pooling}(\hat{\mathbf{G}}_{i,l}^{\text{P}}), \\ \mathbf{f}_{i,l}^{\text{S}} &= \text{max-pooling}(\hat{\mathbf{G}}_{i,l}^{\text{S}}), \\ \mathbf{f}_{i,l}^{\text{PS}} &= \text{relu}(\text{linear}([\mathbf{f}_{i,l}^{\text{P}}]^\top; [\mathbf{f}_{i,l}^{\text{S}}]^\top)). \end{aligned} \quad (11)$$

We concatenate such vectors for all the six pairs of matrices into a vector  $\mathbf{f}_{i,l}$ :

$$\mathbf{f}_{i,l} = [(\mathbf{f}_{i,l}^{\text{PS}})^\top; (\mathbf{f}_{i,l}^{\text{PQ}})^\top; (\mathbf{f}_{i,l}^{\text{PO}})^\top; (\mathbf{f}_{i,l}^{\text{SQ}})^\top; (\mathbf{f}_{i,l}^{\text{SO}})^\top; (\mathbf{f}_{i,l}^{\text{QO}})^\top]^\top. \quad (12)$$

We pack together such vectors for all the  $k'$  paragraphs in  $\mathbf{P}_i^{\text{fus}}$  into a matrix  $\mathbf{F}_i$ :

$$\mathbf{F}_i = [\mathbf{f}_{i,1}; \dots; \mathbf{f}_{i,k'}]. \quad (13)$$

Then we perform *inter-paragraph fusion* over all the  $k'$  paragraphs in  $\mathbf{P}_i^{\text{fus}}$ . We fuse  $\mathbf{F}_i$  into a vector  $\mathbf{f}_i$  by self-attention (Zhong et al., 2019):

$$\begin{aligned} \mathbf{f}_i &= \mathbf{F}_i \times \text{softmax}([a_{i,1}; \dots; a_{i,k'}]^\top), \text{ where} \\ a_{i,l} &= \tanh(\text{linear}(\tanh(\text{linear}(\mathbf{f}_{i,l}))))). \end{aligned} \quad (14)$$

**Option Scoring.** We leverage both the fused dense representation  $\mathbf{f}_i$  and separate paragraph scores  $\mathbf{z}_i^{\text{den}}$  to compute two scores for option  $O_i$ :

$$s_i^{\text{fus}} = \text{linear}(\mathbf{f}_i), \quad s_i^{\text{den}} = \sum_{p_l \in \mathbf{P}_i^{\text{top}}} z_{i,l}^{\text{den}}. \quad (15)$$

**Loss Function.** Given a labeled correct answer  $O_j \in \mathbf{O}$ , we calculate the cross-entropy loss:

$$\begin{aligned} \mathcal{L}^{\text{READ}} &= -\log \frac{\exp(s_j^{\text{fus}})}{\sum_{O_i \in \mathbf{O}} \exp(s_i^{\text{fus}})} - \log \frac{\exp(s_j^{\text{den}})}{\sum_{O_i \in \mathbf{O}} \exp(s_i^{\text{den}})} \\ &= -\log \frac{\exp(s_j^{\text{fus}} + s_j^{\text{den}})}{\sum_{O_i \in \mathbf{O}} \sum_{O_{i'} \in \mathbf{O}} \exp(s_i^{\text{fus}} + s_{i'}^{\text{den}})}. \end{aligned} \quad (16)$$

### 3.5 Joint Training and Prediction

**Joint Training.** We sum the retriever’s loss and the reader’s loss:

$$\mathcal{L} = \mathcal{L}^{\text{RETR}} + \mathcal{L}^{\text{READ}}. \quad (17)$$

We sum such losses over all the training questions.

**Prediction.** For each option  $O_i \in \mathbf{O}$ , we have computed three scores:  $s_i^{\text{spa}}$ ,  $s_i^{\text{den}}$ , and  $s_i^{\text{fus}}$ . We normalize each score over all the  $m$  options by softmax and then calculate their weighted sum as the final score  $s_i$  for  $O_i$ :

$$\begin{aligned} [\hat{s}_1^{\text{spa}}; \dots; \hat{s}_m^{\text{spa}}]^\top &= \text{softmax}([s_1^{\text{spa}}; \dots; s_m^{\text{spa}}]^\top), \\ [\hat{s}_1^{\text{den}}; \dots; \hat{s}_m^{\text{den}}]^\top &= \text{softmax}([s_1^{\text{den}}; \dots; s_m^{\text{den}}]^\top), \\ [\hat{s}_1^{\text{fus}}; \dots; \hat{s}_m^{\text{fus}}]^\top &= \text{softmax}([s_1^{\text{fus}}; \dots; s_m^{\text{fus}}]^\top), \\ s_i &= \alpha \cdot \hat{s}_i^{\text{spa}} + \beta \cdot \hat{s}_i^{\text{den}} + \gamma \cdot \hat{s}_i^{\text{fus}}, \end{aligned} \quad (18)$$

where  $\alpha, \beta, \gamma$  are hyperparameters. We output the option with the highest final score as the answer.

## 4 Experiments

All the experiments were performed on RTX 3090.

### 4.1 Datasets

To our knowledge, there were two datasets containing scenario-based multiple-choice ( $m = 4$ ) questions, both collected from China’s high-school exams. We constructed a new dataset. As shown in Table 1, scenarios and options are paragraph-long and sentence-long, respectively, thereby challenging existing retrievers as discussed in Section 1.

**GeoSQA** (Huang et al., 2019) contains geography questions. Each scenario contains a paragraph and a diagram annotated with a description of its main content. We filtered out questions where options are identifiers for objects in diagrams since diagram understanding is outside our research.

**GH577** (Cheng et al., 2016) contains history questions. Each scenario contains a paragraph.

	GeoSQA	GKMC	GH577
Questions	3,910	1,600	577
Chinese characters per scenario	153	82	41
Chinese characters per option	10	11	15

Table 1: Dataset statistics.

**GKMC** standing for **GaoKao-level multiple-choice** questions is a new dataset we constructed containing geography questions. We followed the procedure in Huang et al. (2019) to crawl and deduplicate questions, but we only kept diagram-free questions to be disjoint from GeoSQA.

**Train-Dev-Test Splits.** We performed five-fold cross-validation with a 3:1:1 split of each dataset into training, development, and test sets.

### 4.2 Corpora

For GeoSQA and GKMC, we combined three geography textbooks and the Chinese Wikipedia into a corpus. We filtered Wikipedia by only keeping paragraphs in the geography domain identified by a BERT-based classifier fine-tuned with 500 positive and 500 negative paragraphs we manually annotated. The corpus contains 175k paragraphs.

For GH577 we used all the 5.6m paragraphs in the Chinese Wikipedia as a corpus.

### 4.3 Implementation Details

JEEVES builds on a PLM. We implemented two variants of JEEVES based on different PLMs for Chinese: BERT-wwm-ext (Cui et al., 2019) and ERNIE (Sun et al., 2019). We coarse-tuned PLMs on C3 (<https://dataset.org/c3/>), a Chinese multiple-choice MRC dataset. We used the BertAdam optimizer. We set maximum sequence length = 256, hidden layer = 12, hidden units = 768, attention heads = 12, dropout rate = 0.1, batch size = 16, learning rate =  $1e-5$ , and warm-up proportion = 0.1. For GeoSQA and GKMC we set epochs = 6 and used two random seeds  $\{1, 2\}$ . For GH577 we set epochs = 10 and used four random seeds  $\{1, 2, 3, 4\}$ .

We used Lucene to implement our retriever. We equipped Lucene’s searcher with a custom scorer using word weights  $w_i$  to retrieve  $\mathbf{P}_i^{\text{top}}$ .

We tuned six main hyperparameters of JEEVES:

- $k \in \{5, 10, 15\}$  below Eq. (5), the number of retrieved paragraphs  $\mathbf{P}_i^{\text{top}}$ ;
- $k' \in \{2, 3, 4\}$  below Eq. (10), the number of fused paragraphs  $\mathbf{P}_i^{\text{fus}}$ ;

- $\tau \in \{100, 200, 300\}$  below Eq. (6), the number of values in  $\mathbf{z}_i^{\text{spa}}$  used for computing  $s_i^{\text{spa}}$ ;
- $\alpha, \beta, \gamma \in \{0, 0.05, \dots, 1\}$  in Eq. (18), the weights of the three scores for each option.

#### 4.4 Baselines

We compared with a variety of baseline methods.

We compared with two retriever-reader methods: **AR** (Pirtoaca et al., 2019) and **DPR** (Karpukhin et al., 2020). We adapted the reader in DPR to multiple-choice questions. We coarse-tuned the retrievers in AR (i.e., document relevance discriminator) and DPR on DuReader (<https://github.com/baidu/DuReader>), a Chinese ODQA dataset containing relevance labels of paragraphs, and we coarse-tuned the readers in AR (i.e., answer verifier discriminator) and DPR on C3. We implemented the following variants of DPR using different strategies for distant supervision.

- **DPR<sub>exact</sub>** represents the original implementation in Karpukhin et al. (2020). Paragraphs that contain an exact match of the gold-standard answer were labeled as positives.
- **DPR<sub>appr- $r$</sub>**  adopts approximate matching. Paragraphs that contain  $\geq r\%$  of the non-stopwords in the gold-standard answer were labeled as positives. We set  $r \in \{100, 50\}$ .
- **DPR<sub>silver</sub>** adopts an integer linear programming based method for distant supervision (Wang et al., 2019) to label  $L$  paragraphs as positives. We set  $L = 3$  according to the statistics of the development sets.

Table 2 shows the proportions of questions where  $\geq 1$  positive paragraph was labeled. Observe that **DPR<sub>exact</sub>** generated positives for 22.28–40.69% of the questions in SQA datasets, while it generated positives for 76.68–83.15% of the questions in ODQA datasets (Karpukhin et al., 2020), showing the difficulty of distant supervision for SQA. So we added **DPR<sub>appr- $r$</sub>**  and **DPR<sub>silver</sub>** which generated more positives but might generate false positives.

We compared with a retrieval-based method: **IR Solver** (Clark et al., 2016). It concatenated the scenario, question, and each option into a query to retrieve the top-ranked paragraph, and used the score of that paragraph as the score of the option.

We compared with five MRC methods: **SPC**, **MMM** (Jin et al., 2020), **DCMN+** (Zhang et al., 2020a), **SG-NET** (Zhang et al., 2020b), and

	GeoSQA	GKMC	GH577
<b>DPR<sub>exact</sub></b>	35.68%	40.69%	22.28%
<b>DPR<sub>appr-100</sub></b>	52.43%	61.06%	33.16%
<b>DPR<sub>appr-50</sub></b>	90.18%	94.00%	79.65%
<b>DPR<sub>silver</sub></b>	100.00%	100.00%	100.00%

Table 2: Questions with positive paragraphs labeled by different strategies for distant supervision.

**DHC** (Liu et al., 2020b). We fed them with top-10 paragraphs relevant to each option. We retrieved such paragraphs in a standard way: by concatenating the scenario, question, and option into a query to retrieve using BM25. **SPC** standing for sentence pair classification employed a sentence pair classifier based on a PLM to rank options by pairing each option with a concatenation of the scenario and the question. For **MMM** we coarse-tuned its sentence encoder using **CMNLI** (<https://github.com/CLUEbenchmark/CLUE>), a Chinese natural language inference dataset, and we used **C3** as its in-domain source dataset for fine-tuning. For **SG-NET** we used **HanLP** (<https://github.com/hankcs/HanLP>) as its parser.

**AR**, **DPR**, **SPC**, **MMM**, **DCMN+**, **SG-NET**, and **DHC** build on a PLM. For a fair comparison with **JEEVES**, we implemented two variants of these methods based on **BERT-wwm-ext** and **ERNIE**, configured in the same way as described in Section 4.3. For readers we tuned maximum sequence length in  $\{256, 512\}$ , epochs in  $\{6, 10\}$ , and learning rate in  $\{1e-5, 3e-5\}$ . For the retriever in **DPR** we tuned maximum sequence length in  $\{128, 256\}$ , set epochs = 40, and other hyperparameters according to Karpukhin et al. (2020).

More implementation details about the baseline methods are on GitHub.

#### 4.5 Evaluation Metrics

To evaluate SQA, we relied on QA labels to measure the **accuracy** of each approach: the proportion of correctly answered questions.

To evaluate retrievers for SQA, we sampled 200 questions randomly but evenly from the test set of each fold on GKMC, and for each question we manually annotated relevant paragraphs in the corpus. We measured the relevance of the paragraphs retrieved by each retriever using **MAP** and **NDCG** at different positions (@2 and @10) in the ranked list of retrieved paragraphs. We also measured **Hit Rate**: the proportion of questions for which  $\geq 1$  retrieved paragraph is relevant.

	GeoSQA		GKMC		GH577	
	dev	test	dev	test	dev	test
IR Solver	31.71 *	31.71 *	45.94 *	45.94 *	33.33 *	33.33 *
BERT-wwm-ext						
AR	34.78 *	33.40 *	51.06 *	50.19 *	34.21 *	31.75 *
DPR <sub>exact</sub>	31.20 *	29.82 *	40.50 *	38.86 *	34.85 *	32.84 *
DPR <sub>appr-100</sub>	32.16 *	30.62 *	41.08 *	39.12 *	35.21 *	32.84 *
DPR <sub>appr-50</sub>	31.68 *	29.87 *	41.72 *	41.66 *	39.41 *	32.29 *
DPR <sub>silver</sub>	31.46 *	31.33 *	40.42 *	39.06 *	35.03 *	34.85 *
SPC	36.39 *	33.84 *	49.81 *	48.25 *	31.58 *	28.07 *
MMM	35.96 *	32.86 *	50.68 *	47.69 *	40.53 *	33.68 *
DCMN+	36.19 *	34.58 *	50.56 *	48.38 *	37.72 *	33.51 *
SG-Net	33.76 *	33.73 *	50.94 *	47.69 *	40.35 *	33.86 *
DHC	35.24 *	33.84 *	48.38 *	48.25 *	40.35 *	32.81 *
JEEVES	<b>38.81</b>	<b>36.95</b>	<b>57.34</b>	<b>55.78</b>	<b>46.27</b>	<b>40.35</b>
ERNIE						
AR	34.45 *	33.30 *	49.63 *	47.31 *	34.39 *	32.81 *
DPR <sub>exact</sub>	32.01 *	29.95 *	41.70 *	39.90 *	36.14 *	32.46 *
DPR <sub>appr-100</sub>	30.75 *	30.99 *	40.37 *	38.06 *	36.14 *	28.24 *
DPR <sub>appr-50</sub>	31.69 *	31.25 *	42.50 *	40.25 *	35.26 *	31.23 *
DPR <sub>silver</sub>	32.32 *	32.20 *	41.56 *	40.19 *	36.84 *	32.51 *
SPC	33.43 *	32.35 *	47.63 *	46.06 *	36.67 *	31.76 *
MMM	34.50 *	31.66 *	45.88 *	40.56 *	36.49 *	31.23 *
DCMN+	34.12 *	32.33 *	49.06 *	47.19 *	36.14 *	31.58 *
SG-Net	32.40 *	31.46 *	48.57 *	46.56 *	37.89 *	34.21 *
DHC	32.84 *	33.09 *	46.69 *	47.31 *	36.49 *	27.02 *
JEEVES	<b>38.76</b>	<b>35.99</b>	<b>57.40</b>	<b>53.78</b>	<b>42.11</b>	<b>36.14</b>

Table 3: Comparison with baselines (accuracy of SQA). We mark the results of baselines that are significantly lower than JEEVES under  $p < 0.01$  (\*).

	BERT-wwm-ext	ERNIE
SPC	102m	99m
MMM	109m	106m
DCMN+	109m	106m
SG-Net	135m	133m
DHC	140m	137m
JEEVES	130m	127m

Table 4: Comparison with baselines (model size).

## 4.6 Results

We averaged the results over all the random seeds and folds on each dataset.

**Comparison with Baselines.** In Table 3, JEEVES achieved the highest accuracy of SQA on both the validation and test sets of all the three datasets. It significantly outperformed all the baselines under  $p < 0.01$ . The strongest baseline methods were DCMN+ (BERT-wwm-ext) on the test set of GeoSQA, AR (BERT-wwm-ext) on GKMC, and SG-Net (ERNIE) on GH577. JEEVES (BERT-wwm-ext) exceeded them by 2.37 of accuracy on GeoSQA, by 5.59 on GKMC, and by 6.14 on GH577. The results demonstrated the effectiveness of our end-to-end model for SQA.

Observe that JEEVES achieved better performance by a model comparable to baseline methods in size. As shown in Table 4, their numbers of parameters were of the same order of magnitude.

**Comparison between Retrievers.** In Table 5, JEEVES retrieved the most relevant paragraphs for SQA. It significantly outperformed all the baseline retrievers under  $p < 0.01$ . In particular, JEEVES (BERT-wwm-ext) largely exceeded

	Hit Rate		MAP		NDCG	
	@2	@10	@2	@10	@2	@10
BM25	32.76 *	56.57 *	29.43 *	34.11 *	24.43 *	31.02 *
IR Solver	37.72 *	62.00 *	34.00 *	38.44 *	29.54 *	35.39 *
BERT-wwm-ext						
AR	37.62 *	56.86 *	31.95 *	35.23 *	27.03 *	32.13 *
DPR <sub>exact</sub>	25.01 *	45.57 *	21.14 *	23.50 *	16.15 *	16.20 *
DPR <sub>appr-100</sub>	27.65 *	48.38 *	23.69 *	26.20 *	17.92 *	17.49 *
DPR <sub>appr-50</sub>	30.80 *	53.33 *	26.35 *	29.04 *	20.19 *	20.20 *
DPR <sub>silver</sub>	25.64 *	44.86 *	21.43 *	23.70 *	16.13 *	15.67 *
JEEVES	<b>42.90</b>	<b>67.90</b>	<b>37.00</b>	<b>39.47</b>	<b>31.31</b>	<b>37.91</b>
ERNIE						
AR	38.85 *	60.00 *	34.47 *	37.39 *	29.11 *	33.66 *
DPR <sub>exact</sub>	25.46 *	43.94 *	21.72 *	24.06 *	16.40 *	15.78 *
DPR <sub>appr-100</sub>	25.34 *	45.66 *	21.91 *	24.25 *	16.43 *	16.25 *
DPR <sub>appr-50</sub>	30.27 *	49.69 *	26.02 *	27.95 *	19.68 *	18.61 *
DPR <sub>silver</sub>	24.14 *	44.50 *	20.85 *	23.51 *	15.41 *	14.90 *
JEEVES	<b>41.33</b>	<b>65.43</b>	<b>36.83</b>	<b>39.22</b>	<b>31.19</b>	<b>36.94</b>

Table 5: Comparison between retrievers on the test set of GKMC. We mark the results of baselines that are significantly lower than JEEVES under  $p < 0.01$  (\*).

	GeoSQA		GKMC		GH577	
	dev	test	dev	test	dev	test
BERT-wwm-ext						
SPC	32.30 *	31.46 *	37.66 *	36.83 *	37.37 *	31.93 *
MMM	35.81 *	35.24 °	52.00 *	50.88 *	40.70 *	35.09 *
DCMN+	35.22 *	33.17 *	50.44 *	46.94 *	40.35 *	37.19 °
SG-Net	34.99 *	34.48 *	48.25 *	47.06 *	39.12 *	31.05 *
DHC	34.94 *	33.58 *	49.13 *	48.44 *	43.33 °	38.77 *
JEEVES	<b>38.81</b>	<b>36.95</b>	<b>57.34</b>	<b>55.78</b>	<b>46.27</b>	<b>40.35</b>
ERNIE						
SPC	33.35 *	32.22 *	39.31 *	36.00 *	39.30 *	31.93 *
MMM	32.51 *	32.45 *	43.69 *	41.31 *	35.44 *	29.47 *
DCMN+	33.63 *	32.71 *	45.19 *	43.31 *	38.60 *	32.28 *
SG-Net	33.07 *	32.20 *	47.13 *	44.50 *	38.42 *	33.68 *
DHC	33.32 *	32.15 *	45.19 *	44.75 *	35.79 *	34.56 *
JEEVES	<b>38.76</b>	<b>35.99</b>	<b>57.40</b>	<b>53.78</b>	<b>42.11</b>	<b>36.14</b>

Table 6: Comparison between readers (accuracy of SQA). We mark the results of baselines that are significantly lower than JEEVES under  $p < 0.01$  (\*) or  $p < 0.05$  (°).

BM25 by 10.14–11.33 of Hit Rate, 5.36–7.57 of MAP, and 6.88–6.89 of NDCG. Among supervised baseline retrievers, DPR<sub>appr-50</sub> was the best-performing variant of DPR but still inferior to AR. Interestingly, none of them were comparable with the unsupervised IR Solver, showing the limitations of existing distant supervision (e.g., DPR) and transfer learning (e.g., AR) based retrievers for SQA. The results demonstrated the effectiveness of our implicitly supervised retriever for SQA.

**Comparison between Readers.** To compare our reader with baseline readers, we fed the five MRC methods with paragraphs retrieved by our retriever. In Table 6, JEEVES remained to achieve the highest accuracy of SQA on both the validation and test sets of all the three datasets. It significantly outperformed all the baseline readers under  $p < 0.01$  or  $p < 0.05$  except for DHC on the test set of GH577 where the difference was not significant. In particular, JEEVES (BERT-wwm-ext) noticeably exceeded the best-performing baseline reader on the

	GeoSQA		GKMC		GH577	
	dev	test	dev	test	dev	test
BERT-wwm-ext						
JEEVES	<b>38.81</b>	<b>36.95</b>	<b>57.34</b>	<b>55.78</b>	<b>46.27</b>	<b>40.35</b>
w/o $s_i^{\text{den}}$	37.16 •	34.44 •	54.28 •	50.53 •	42.72 •	36.27 •
w/o $s_i^{\text{fus}}$	37.28 •	35.10 •	52.75 •	49.88 •	45.00 °	38.03
w/o joint training	37.24 •	36.24	52.59 •	46.96 •	45.09 °	38.60
ERNIE						
JEEVES	<b>38.76</b>	<b>35.99</b>	<b>57.40</b>	<b>53.78</b>	<b>42.11</b>	<b>36.14</b>
w/o $s_i^{\text{den}}$	37.28 °	34.30 •	53.56 •	50.47 •	39.39 °	33.55 °
w/o $s_i^{\text{fus}}$	36.89 •	35.23	53.69 •	50.94 •	39.52 °	34.39
w/o joint training	36.02 •	33.68 •	51.65 •	47.72 •	41.92	34.21

Table 7: Ablation study (accuracy of SQA). We mark the results of reduced versions of JEEVES that are significantly lower than the full version of JEEVES under  $p < 0.01$  (•) or  $p < 0.05$  (°).

$k$	dev	test	$k'$	dev	test	$\tau$	dev	test
5	55.56	52.44	2	<b>57.34</b>	<b>55.78</b>	100	56.75	54.19
10	<b>57.34</b>	<b>55.78</b>	3	57.00	54.62	200	<b>57.34</b>	<b>55.78</b>
15	56.69	54.19	4	57.19	53.38	300	55.37	53.75

(a)  $k$                       (b)  $k'$                       (c)  $\tau$

Table 8: Hyperparameter study (accuracy of SQA) of JEEVES (BERT-wwm-ext) on GKMC.

test set of each dataset, i.e., MMM (BERT-wwm-ext) on GeoSQA and GKMC, and DHC (BERT-wwm-ext) on GH577, by 1.71, 4.90, and 1.58 of accuracy, respectively. The results demonstrated the effectiveness of our reader for SQA.

**Ablation Study.** In JEEVES we computed three scores for each option:  $s_i^{\text{spa}}$  in Eq. (6) based on sparse representations of paragraphs,  $s_i^{\text{den}}$  in Eq. (15) based on dense representations of separate paragraphs, and  $s_i^{\text{fus}}$  in Eq. (15) based on fused dense representations of paragraphs. For ablation study we could not remove  $s_i^{\text{spa}}$  because we relied on it to implicitly supervise our retriever. We removed  $s_i^{\text{den}}$  and  $s_i^{\text{fus}}$  to assess their usefulness. In Table 7, compared with the full version of JEEVES, the two reduced versions both fell in accuracy on both the validation and test sets of all the three datasets. By removing  $s_i^{\text{den}}$ , the accuracy decreased significantly under  $p < 0.01$  or  $p < 0.05$  in all the cases. By removing  $s_i^{\text{fus}}$ , the decreases were also significant in most cases. The results demonstrated the effectiveness of using both separate and fused dense representations of paragraphs for SQA.

In JEEVES we jointly trained our retriever and reader. For ablation study we separately trained them. In Table 7, compared with the full version of JEEVES, the reduced version fell in accuracy on both the validation and test sets of all the three datasets. Without joint training, the accuracy decreased significantly under  $p < 0.01$  or  $p < 0.05$  in most cases. The results demonstrated the effectiveness of joint training for SQA.

Source of Error	GeoSQA	GKMC	GH577
Corpus	88%	22%	14%
Retriever	8%	36%	48%
Reader	14%	66%	64%

Table 9: Error analysis of JEEVES (BERT-wwm-ext).

**Hyperparameter Study.** We reported the accuracy of JEEVES (BERT-wwm-ext) on GKMC under different values of its three main hyperparameters:  $k \in \{5, 10, 15\}$ ,  $k' \in \{2, 3, 4\}$ , and  $\tau \in \{100, 200, 300\}$ . In Table 8, the highest accuracy was observed on the validation set under  $k = 10$ ,  $k' = 2$ , and  $\tau = 200$ . Therefore, we used these settings throughout all the above experiments.

**Error Analysis.** We randomly sampled 50 questions from the test set of each dataset to which JEEVES (BERT-wwm-ext) outputted an incorrect answer. We analyzed the source of each error. Note that an error could have multiple sources. In Table 9, corpus was the main source of error on GeoSQA. In particular, commonsense knowledge was needed but often absent in our corpus. Retrieval errors were frequent on GKMC and GH577. Our retriever sometimes assigned large weights to noise words in scenarios and retrieved irrelevant paragraphs. Besides, it could not identify some semantic matches such as *haze* and *air pollution*. Most errors on GKMC and GH577 were related to our reader. Indeed, it could not offer advanced reasoning capabilities such as comparison and negation, which were needed for some questions.

**Run Time.** JEEVES used an average of 0.2 second for answering a question in the test sets.

## 5 Conclusion

Scenario-based multiple-choice questions have posed a great challenge to the retriever-reader framework: keyphrases for retrieving relevant paragraphs are blended with much noise in a long scenario description. In the absence of relevance labels of paragraphs, we devised a novel word weighting mechanism to implicitly train our retriever only using QA labels. It significantly outperformed existing unsupervised, distant supervision based, and transfer learning based retrievers. Based on that, our joint end-to-end model JEEVES exceeded a variety of strong baseline methods on three SQA datasets. While our experiments in the paper are focused on multiple-choice questions in high-school exams, JEEVES has the potential to be adapted

to other SQA tasks where scenarios are also very long, such as legal judgment predication. It would be interesting to explore this direction in the future.

Error analysis has revealed some shortcomings of JEEVES, which we will address in future work. Among others, we will incorporate knowledge graphs (Shen et al., 2021; Zhang et al., 2018; Li et al., 2020) and enhance reasoning capabilities (Sun et al., 2018).

## Acknowledgments

This work was supported by the National Key R&D Program of China (2018YFB1005100).

## References

- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. [Neural legal judgment prediction in English](#). In *ACL*, pages 4317–4323.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to answer open-domain questions](#). In *ACL*, pages 1870–1879.
- Gong Cheng, Weixi Zhu, Ziwei Wang, Jianghui Chen, and Yuzhong Qu. 2016. Taking up the Gaokao challenge: An information retrieval approach. In *IJCAI*, pages 2479–2485.
- Peter Clark, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter D. Turney, and Daniel Khashabi. 2016. Combining retrieval, statistics, and inference to answer elementary science questions. In *AAAI*, pages 2580–2586.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Ziqing Yang, Shijin Wang, and Guoping Hu. 2019. Pre-training with whole word masking for Chinese BERT. *CoRR*, abs/1906.08101.
- Zixian Huang, Yulin Shen, Xiao Li, Yuang Wei, Gong Cheng, Lin Zhou, Xinyu Dai, and Yuzhong Qu. 2019. [GeoSQA: A benchmark for scenario-based question answering in the geography domain at high school level](#). In *EMNLP-IJCNLP*, pages 5865–5870.
- Di Jin, Shuyang Gao, Jiun-Yu Kao, Tagyoung Chung, and Dilek Hakkani-Tür. 2020. MMM: multi-stage multi-task learning for multi-choice reading comprehension. In *AAAI-IAAI-EAAI*, pages 8010–8017.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *EMNLP*, pages 6769–6781.
- Omar Khattab and Matei Zaharia. 2020. [Colbert: Efficient and effective passage search via contextualized late interaction over BERT](#). In *SIGIR*, pages 39–48.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural Questions: a benchmark for question answering research. *Trans. Assoc. Comput. Linguistics*, 7:452–466.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard H. Hovy. 2017. [RACE: large-scale reading comprehension dataset from examinations](#). In *EMNLP*, pages 785–794.
- Adam Lally, Sugato Bagchi, Michael Barborak, David W. Buchanan, Jennifer Chu-Carroll, David A. Ferrucci, Michael R. Glass, Aditya Kalyanpur, Erik T. Mueller, J. William Murdock, Siddharth Patwardhan, and John M. Prager. 2017. [Watson-Paths: Scenario-based question answering and inference over unstructured information](#). *AI Magazine*, 38(2):59–76.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *ACL*, pages 6086–6096.
- Shuxin Li, Zixian Huang, Gong Cheng, Evgeny Kharlamov, and Kalpa Gunaratna. 2020. [Enriching documents with compact, representative, relevant knowledge graphs](#). In *IJCAI*, pages 1748–1754.
- Xiao Li, Yawei Sun, and Gong Cheng. 2021. [TSQA: tabular scenario based question answering](#). In *AAAI*.
- Dayiheng Liu, Yeyun Gong, Jie Fu, Yu Yan, Jiusheng Chen, Daxin Jiang, Jiancheng Lv, and Nan Duan. 2020a. [RikiNet: reading Wikipedia pages for natural question answering](#). In *ACL*, pages 6762–6771.
- Zhuang Liu, Kaiyu Huang, Degen Huang, Zhuang Liu, and Jun Zhao. 2020b. [Dual head-wise coattention network for machine comprehension with multiple-choice questions](#). In *CIKM*, pages 1015–1024.
- Yoshitomo Matsubara, Thuy Vu, and Alessandro Moschitti. 2020. [Reranking for efficient Transformer-based answer selection](#). In *SIGIR*, pages 1577–1580.
- George-Sebastian Pirtoaca, Traian Rebedea, and Stefan Ruseti. 2019. [Answering questions by learning to rank - learning to rank by answering questions](#). In *EMNLP-IJCNLP*, pages 2531–2540.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100, 000+ questions for machine comprehension of text](#). In *EMNLP*, pages 2383–2392.
- Yulin Shen, Ziheng Chen, Gong Cheng, and Yuzhong Qu. 2021. [CKGG: a Chinese knowledge graph for high-school geography education and beyond](#). In *ISWC*.

- Yawei Sun, Gong Cheng, and Yuzhong Qu. 2018. Reading comprehension with graph-based temporal-casual reasoning. In *COLING*, pages 806–817.
- Yu Sun, Shuohuan Wang, Yu-Kun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. ERNIE: enhanced representation through knowledge integration. *CoRR*, abs/1904.09223.
- Hai Wang, Dian Yu, Kai Sun, Jianshu Chen, Dong Yu, David A. McAllester, and Dan Roth. 2019. Evidence sentence extraction for machine reading comprehension. In *CoNLL*, pages 696–707.
- Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerry Tesauro, Bowen Zhou, and Jing Jiang. 2018. R<sup>3</sup>: Reinforced ranker-reader for open-domain question answering. In *AAAI-IAAI-EAAI*, pages 5981–5988.
- Nuo Xu, Pinghui Wang, Long Chen, Li Pan, Xiaoyan Wang, and Junzhou Zhao. 2020. Distinguish confusing law articles for legal judgment prediction. In *ACL*, pages 3086–3095.
- Wenmian Yang, Weijia Jia, Xiaojie Zhou, and Yutao Luo. 2019. Legal judgment prediction via multi-perspective bi-feedback network. In *IJCAI*, pages 4085–4091.
- Shuailiang Zhang, Hai Zhao, Yuwei Wu, Zhuosheng Zhang, Xi Zhou, and Xiang Zhou. 2020a. DCMN+: dual co-matching network for multi-choice reading comprehension. In *AAAI-IAAI-EAAI*, pages 9563–9570.
- Zhiwei Zhang, Lingling Zhang, Hao Zhang, Weizhuo He, Zequn Sun, Gong Cheng, Qizhi Liu, Xinyu Dai, and Yuzhong Qu. 2018. Towards answering geography questions in Gaokao: A hybrid approach. In *CCKS*, pages 1–13.
- Zhuosheng Zhang, Yuwei Wu, Junru Zhou, Sufeng Duan, Hai Zhao, and Rui Wang. 2020b. SG-Net: syntax-guided machine reading comprehension. In *AAAI-IAAI-EAAI*, pages 9636–9643.
- Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Chaojun Xiao, Zhiyuan Liu, and Maosong Sun. 2018. Legal judgment prediction via topological learning. In *EMNLP*, pages 3540–3549.
- Victor Zhong, Caiming Xiong, Nitish Shirish Keskar, and Richard Socher. 2019. Coarse-grain fine-grain coattention network for multi-evidence question answering. In *ICLR*.