

Joint Learning of Answer Selection and Answer Summary Generation in Community Question Answering*

Yang Deng,¹ Wai Lam,¹ Yuexiang Xie,³ Daoyuan Chen,⁴ Yaliang Li,⁴ Min Yang,⁵ Ying Shen^{2,3†}

[†]Corresponding Author

¹The Chinese University of Hong Kong, ²South China University of Technology,

³Peking University Shenzhen Graduate School, ⁴Alibaba Group, ⁵Chinese Academy of Sciences,

{ydeng, wlam}@se.cuhk.edu.hk, {xieyx, shenyings}@pku.edu.cn,
{daoyuanchen.cdy, yaliang.li}@alibaba-inc.com, min.yang@siat.ac.cn

Abstract

Community question answering (CQA) gains increasing popularity in both academy and industry recently. However, the redundancy and lengthiness issues of crowdsourced answers limit the performance of answer selection and lead to reading difficulties and misunderstandings for community users. To solve these problems, we tackle the tasks of answer selection and answer summary generation in CQA with a novel joint learning model. Specifically, we design a question-driven pointer-generator network, which exploits the correlation information between question-answer pairs to aid in attending the essential information when generating answer summaries. Meanwhile, we leverage the answer summaries to alleviate noise in original lengthy answers when ranking the relevancy degrees of question-answer pairs. In addition, we construct a new large-scale CQA corpus, WikiHowQA, which contains long answers for answer selection as well as reference summaries for answer summarization. The experimental results show that the joint learning method can effectively address the answer redundancy issue in CQA and achieves state-of-the-art results on both answer selection and text summarization tasks. Furthermore, the proposed model is shown to be of great transferring ability and applicability for resource-poor CQA tasks, which lack of reference answer summaries.

Introduction

Recent years have witnessed a spectacular increase in real-world applications of community question answering (CQA), such as Yahoo! Answer¹ and StackExchange². Many studies have been made on different tasks in CQA, such as answer selection, question-question relatedness, and comment classification (Moschitti, Bonadiman, and Uva 2017; Joty, Márquez, and Nakov 2018; Nakov et al. 2017). However, due to the length and redundancy of answers in CQA scenario, there are several challenges that need to be

tackled in real-world applications. (i) The noise introduced by the redundancy of answers makes it difficult for answer selection model to pick out correct answers from a set of candidates. (ii) Compared with other QA systems (e.g., factoid question answering), answers in CQA are often too long for community users to read and comprehend.

Current state-of-the-art answer selection models (Tan et al. 2016; Wu, Sun, and Wang 2018) employ the attention mechanism to attend the important correlated information between question-answer pairs. These methods perform well when ranking short answers, while the accuracy goes down with the increase in the length of answers (dos Santos et al. 2016; Rücklé, Moosavi, and Gurevych 2019). Recent studies on coarse-to-fine question answering for long documents, such as Reading Comprehension (RC) (Choi et al. 2017; Wang et al. 2018a; 2018b), focus on the answer span extraction in factoid QA, in which those factoid questions can be answered by a certain word or a short phrase. Conversely, in non-factoid CQA, discrete and complex information from multiple sentences makes up the answers together. Besides, generative RC methods (Nishida et al. 2019) only give one certain answer, while there are often multiple useful answers in CQA. Thus, these approaches are not suitable for addressing the redundancy issue of answers in CQA.

On the other hand, text summarization provides an effective approach to alleviating the aforementioned issue. Text summarization methods can generally be divided into two categories: extractive summarization (Cheng and Lapata 2016; Nallapati, Zhai, and Zhou 2017) and abstractive summarization (See, Liu, and Manning 2017; Nallapati et al. 2016). The aim is to assemble or generate summaries from the source article or external vocabulary, based on the information from the source text. In the existing studies, answer summarization in CQA is mainly explored by extractive summarization models (Tomasoni and Huang 2010; Song et al. 2017). However, due to the length of answers, extractive methods sometimes fall short of generalization of all the important information in the whole answer and consistency of the core idea. Besides, the correlation information between question and answer, which plays a crucial role in human comprehension, is underutilized by current query-based summarization studies (Nema et al. 2017;

*This work was financially supported by the National Natural Science Foundation of China (No.61602013) and a grant from the Research Grant Council of the Hong Kong Special Administrative Region, China (Project Code: 14204418).

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

¹<https://answers.yahoo.com/>

²<https://stackexchange.com/>

Singh et al. 2018). Therefore, we intend to take advantage of both the contextual information from the source text and the relationship between the question-answer pair to generate abstractive answer summaries in CQA.

We aim to simultaneously tackle the above issues in CQA, including (i) improving the performance of non-factoid answer selection with long answers, (ii) generating abstractive summaries of the answers. We jointly learn answer selection and abstractive summarization to generate answer summaries for CQA. First, we exploit the correlated information between question-answer pairs to improve abstractive answer summarization, which enables the summarizer to generate abstractive summaries related to questions. Then, we measure the relevancy degrees between questions and answer summaries to alleviate the impact of noise from original answers. Besides, since obtaining reference summaries is usually labor-intensive and time-consuming in a new domain, a transfer learning strategy is designed to improve resource-poor CQA tasks with large-scale supervision data.

We summarize our contributions as follows:

1. We jointly learn answer selection and answer summary generation to tackle the lengthiness and redundancy issues of the answer in CQA with a unified model. A novel joint learning framework of answer selection and abstractive summarization (ASAS) is proposed to employ the question information to guide the abstractive summarization, and meanwhile leverage the summaries to reduce noise in answers for precisely measuring the correlation degrees of QA pairs.

2. We construct a new dataset, WikiHowQA, for the task of answer summary generation in CQA, which can be adapted to both answer selection and summarization tasks. Experimental results on WikiHowQA show that the proposed joint learning method outperforms SOTA answer selection methods and meanwhile generates more precise answer summaries than existing summarization methods.

3. To handle resource-poor CQA tasks, we design a transfer learning strategy, which enable those tasks without reference answer summaries to conduct the joint learning with impressive experimental results.

Related Work

Community Question Answering. Answer selection is the core and the most widely-studied problem in community question answering. Recent studies have evolved from feature-based methods (Wang, Ming, and Chua 2009; Wang and Manning 2010) into deep learning models, such as convolutional neural network (CNN) (Severyn and Moschitti 2015) and recurrent neural network (RNN) (Wang and Nyberg 2015). In order to capture the interactive information in QA sentences, various attention mechanisms (Tan et al. 2016; dos Santos et al. 2016) are developed to align the related words between questions and answers. However, the lengthy and redundant answers in CQA scenario may introduce much noise and scatter important information, which causes difficulties in answer selection. Some studies leverage additional information to compensate the imbalance of information between questions and answers, such as user model (Wen et al. 2018; Li et al. 2017), latent topic (Yoon, Shin, and Jung 2018), external knowl-

edge (Shen et al. 2018) or question subject (Wu, Sun, and Wang 2018). Some existing transfer learning studies on CQA focus on cross-domain adaptation (Deng et al. 2018; Yu et al. 2018). In this work, we employ summarization method to reduce noise in the original lengthy answers to improve the answer selection performance in CQA.

Text Summarization. Text summarization techniques are mainly classified into two categories: extractive and abstractive summarization. Extractive approaches regard summarization as a sentence classification (Nallapati, Zhai, and Zhou 2017) or a sequence labeling task (Cheng and Lapata 2016) to select sentences from the article to form the summary, while abstractive approaches usually employ attention-based encoder-decoder models (Nallapati et al. 2016; See, Liu, and Manning 2017) to generate abstractive summaries. Answer summarization in CQA was first introduced by Zhou, Lin, and Hovy (2006) as an application of extractive summarization. After that, studies on answer summarization are still regarded as a separate extractive summarization module in QA pipeline (Tomasoni and Huang 2010; Song et al. 2017). Besides, query-based summarization methods (Nema et al. 2017; Singh et al. 2018) also can be a good solution for this task, however, these approaches are reported to perform worse than answer selection methods on question answering scenario (Saha et al. 2018).

Multi-task Learning. Inspired by the success of multi-task learning in other NLP tasks, several attempts have been made to solve answer selection with different tasks. Moschitti, Bonadiman, and Uva (2017) and Joty, Márquez, and Nakov (2018) enhance answer selection in CQA via multi-task learning with the auxiliary tasks of question-question relatedness and question-comment relatedness. Yang et al. (2019) leverage the question categorization to enhance the question representation learning for CQA. Deng et al. (2019) propose a multi-view attention based multi-task learning model to jointly tackle answer selection and knowledge base question answering tasks. In this work, we jointly learn answer selection and abstractive summarization to select and generate precise answers in CQA.

Method

Problem Definition

We aim to jointly conduct two tasks, answer selection and abstractive summarization, to select and generate concise answers for CQA. Given a question q_i , the goal is to simultaneously select the set of correct answers from a set of candidates $A_i = \{a_i^{(1)}, \dots, a_i^{(j)}\}$ and generate an abstractive summary $\beta_i^{(*)}$ for each selected answer $a_i^{(*)}$.

The dataset D for learning typically contains a set of questions Q with the number of N . For each question $q_i \in Q$, there are M_i candidate answers A_i with the corresponding reference summary $\beta_i^{(j)}$ written by human and the label $y_i^{(j)}$ determining whether $a_i^{(j)}$ can answer q_i .

$$D = \{(q_i, \{(a_i^{(j)}, \beta_i^{(j)}, y_i^{(j)})\}_{j=1}^{M_i})\}_{i=1}^N. \quad (1)$$

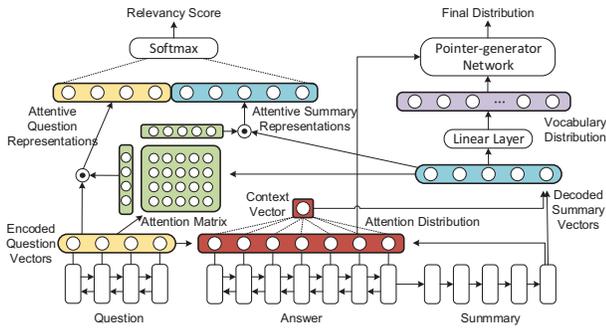


Figure 1: The Joint Learning Framework of Answer Selection and Abstractive Summarization (ASAS).

Model

We introduce the proposed joint learning model for answer selection and abstractive summarization (ASAS). As is depicted in Fig. 1, The overall framework of ASAS consists of four components: (i) Shared Compare-Aggregate Bi-LSTM Encoder, (ii) Sequence-to-sequence Model with Question-aware Attention, (iii) Question Answer Alignment with Summary Representations, (iv) Question-driven Pointer-generator Network.

Shared Compare-Aggregate Bi-LSTM Encoder. The word embeddings of the question and the original answer, W_q and W_a , are fed into a compare layer as Wang and Jiang (2017) to generate the model input \hat{W}_q and \hat{W}_a . Then, a pair of Bi-LSTM encoders are adopted to aggregate the context information. We encode a pair of word sequences of the question q and the answer a into sentence representations $H \in \mathbb{R}^{L \times d_h}$, where L and d_h are the length of sentences and the size of hidden states:

$$H_q = \text{Bi-LSTM}(\hat{W}_q), \quad H_a = \text{Bi-LSTM}(\hat{W}_a). \quad (2)$$

Seq2Seq Model with Question-aware Attention. With the intuition that the information in the question is supposed to be helpful in attending the important elements in the original answer sentence, we propose a question-aware attention based seq2seq model to decode the encoded sentence representation of the answer. We adopt a unidirectional LSTM as the decoder. On each step t , the decoder produces the hidden state s_t with the input of the previous word w_{t-1} . The question-aware attention α^t is generated by:

$$s_t = \text{LSTM}(s_{t-1}, w_{t-1}), \quad (3)$$

$$o_q = \text{Average}(H_q); \quad (4)$$

$$e_i^t = m^t \tanh(W_h h_i^a + W_s s_t + W_q o_q + b), \quad (5)$$

$$\alpha^t = \text{softmax}(e^t), \quad (6)$$

where m , W_h , W_s , W_q are attention parameter matrices to be learned. The question-aware attention weight α^t is used to generate context vector \hat{h}_t as a probability distribution over the source words:

$$\hat{h}_t = \sum_i \alpha^t h_i^a. \quad (7)$$

The context vector aggregates the information from the source text and the question for the current step. We concatenate it with the decoder state s_t and pass through a linear layer to generate the summary representation h_t^s :

$$h_t^s = W_1 [s_t : \hat{h}_t] + b_1, \quad (8)$$

where W_1 and b_1 are parameters to be learned.

Question Answer Alignment with Summary Representations. We apply a two-way attention mechanism to generate the co-attention between the encoded question representation H_q and the decoded summary representation H_s :

$$M_{qa} = \tanh(H_q^T U H_s), \quad (9)$$

$$\alpha_q = \text{softmax}(\text{Max}(M_{qa})), \quad (10)$$

$$\alpha_a = \text{softmax}(\text{Max}(M_{qa}^T)), \quad (11)$$

where $U \in \mathbb{R}^{d_s \times d_s}$ is the attention parameter matrix to be learned; d_s is the dimension of QA representations; α_q and α_a are the co-attention weights for the question and the answer summary respectively.

We conduct dot product between the attention vectors and the question and summary representations to generate the final attentive sentence representations for answer selection:

$$r_q = H_q^T \alpha_q, \quad r_a = H_s^T \alpha_a. \quad (12)$$

Compared with encoded answer representations, decoded summary representations are more concise and compressive, which enable answer selection model to precisely capture the interactive information between questions and answers.

Question-driven Pointer-generator Network. First, the probability distribution P_{vocab} over the fixed vocabulary is obtained by passing the summary representation h_t^s through a softmax layer:

$$P_{vocab} = \text{softmax}(W_2 h_t^s + b_2), \quad (13)$$

where W_2 and b_2 are parameters to be learned. Then, a question-aware pointer network is proposed to copy words from the source article with the guidance of the question information. The question-aware generation probability $p_{gen} \in [0, 1]$ takes into account the decoded summary representation h_t^s , the decoder input x_t and the question representation o_q :

$$p_{gen} = \sigma(w_h^T h_t^s + w_x^T x_t + w_q^T o_q + b_p), \quad (14)$$

where w_h , w_x , w_q and b_p are parameters to be learned, and σ is the sigmoid function. Following the basic pointer-generator network (PGN) (See, Liu, and Manning 2017), we obtain the final probability distribution over both the fixed vocabulary and words from the source article:

$$P = p_{gen} P_{vocab} + (1 - p_{gen}) \sum_{i:w_i=w} \alpha_i^t. \quad (15)$$

To be specific, the question information is involved in not only the generating process, but also the copying process in the question-driven PGN. (i) The question information directs the calculation of the generation probability to decide

whether generating a word from the vocabulary or copying from the source text. (ii) The question-aware attention weights integrate the question information to attend the important words in the source text for copying. (iii) The probability distribution over the vocabulary is learned from the question-aware attentive summary representations.

Joint Training Procedure

Answer Selection Loss. The attentive representations of questions and summaries go through a softmax layer for binary classification:

$$y(q, a) = \text{softmax}(W_s[r_q : r_a] + b_s), \quad (16)$$

where $W_s \in \mathbb{R}^{d_x \times 2}$ and $b_s \in \mathbb{R}^2$ are parameters to be learned. The answer selection task is trained to minimize the cross-entropy loss function:

$$L_{qa} = - \sum_{i=1}^N [y_i \log p_i + (1 - y_i) \log (1 - p_i)], \quad (17)$$

where p is the output of the softmax layer and y is the binary classification label of the QA pair.

Summarization Loss. The summarization task is trained to minimize the negative log likelihood:

$$L_{sum} = - \frac{1}{T} \sum_{t=0}^T \log P(w_t^*). \quad (18)$$

Coverage Loss. Coverage loss (See, Liu, and Manning 2017) was proposed to discourage the repetition in abstractive summarization. In each decoder timestep t , the coverage vector $c^t = \sum_{t'=0}^{t-1} a^{t'}$ is used to represent the degree of coverage so far. The coverage vector c^t will be applied to compute the attention weight α^t . The coverage loss is trained to penalize the repetition in updated attention weight α^t :

$$L_{cov} = \frac{1}{T} \sum_{t=1}^T \sum_i \min(\alpha_i^t, c_i^t). \quad (19)$$

Overall Loss Function. For joint training, the final objective function is to minimize above three loss functions:

$$L = \lambda_1 L_{qa} + \lambda_2 L_{sum} + \lambda_3 L_{cov}, \quad (20)$$

where $\lambda_1, \lambda_2, \lambda_3$ are hyper-parameters to balance losses.

Handling Resource-poor Datasets

Since annotating gold answer summaries is a labor-intensive work, we intend to leverage the knowledge learned from the joint learning of answer selection and answer summary generation on a large-scale supervision dataset and apply it to resource-poor datasets without reference answer summaries. The goal can be achieved by a transfer learning strategy involving two steps: (i) initialize the parameters of model pre-trained on the source dataset, (ii) further fine-tune on the target dataset. A straightforward way is to fine-tune all the parameters learned from the source data on the target training dataset. Another fashion is to fine-tune a certain part of

	train / dev / test
#Questions	76,687 / 8,000 / 22,354
#QA Pairs	904,460 / 72,474 / 211,255
#Summaries	142,063 / 18,909 / 42,624
Avg QLen	7.20 / 6.84 / 6.69
Avg ALen	520.87 / 548.26 / 554.66
Avg SLen	67.38 / 61.84 / 74.42
Avg #CandA	11.79 / 9.06 / 9.45

Table 1: Statistic of WikiHowQA Dataset

parameters and keep the remaining part of model fixed during fine-tuning. In this case, we first pre-train the whole joint learning model on the source dataset, and then only fine-tune the answer selection modules (including Shared Compare-Aggregate Bi-LSTM Encoder & Question Answer Alignment). On one hand, fixing the summarization part can not only reduce the demand for annotating summary data, but also prevent model over-fitting. On the other hand, questioning styles and answer contents vary from CQA tasks in different domains, thus, the answer selection part is supposed to benefit from fine-tuning in target domains.

Datasets and Experimental Setting

Datasets

Most of the widely-adopted answer selection benchmark datasets are composed of short sentences, such as WikiQA (Yang, Yih, and Meek 2015), SemEval (Nakov et al. 2017), WikiPassageQA (Cohen, Yang, and Croft 2018) and StackExchange (Rücklé, Moosavi, and Gurevych 2019), two latest non-factoid answer selection datasets with long passages (about 150 words) as candidate answers, lack of the reference summary for answer summarization evaluation in our defined answer summary generation task.

We present a new CQA corpus, WikiHowQA, for answer summary generation, which contains labels for the answer selection task as well as reference summaries for the text summarization task. To prepare this dataset, we modify a latest text summarization dataset, WikiHow (Koupaei and Wang 2018), which was obtained from *WikiHow*³ knowledge base. The WikiHow dataset contains detailed answers written by community users for non-factoid questions starting with “**How to**”. The original answers are composed by multiple steps of different methods for the question, and the description in each step is associated with an abstractive summary. The WikiHow dataset only contains the selected ground-truth answers and the reference summaries for each answer, while the whole candidate answer set is required when we wish to conduct answer selection experiments on this dataset. Therefore, we construct a new CQA dataset based on the WikiHow dataset.

We first clean up the WikiHow dataset by filtering out those questions without answers or summaries and those answers with punctuation only. After that, the dataset size is reduced from 230,843 to 203,596, including 107,041 unique

³<http://www.wikihow.com/>

StackExchange	#Questions (train/dev/test)	Avg ALen
Travel	3,572 / 765 / 766	214
Cooking	3,692 / 791 / 792	189
Academia	2,856 / 612 / 612	229
Apple	5,831 / 1,249 / 1,250	114
Aviation	3,035 / 650 / 652	281

Table 2: Statistic of StackExchange CQA Dataset

questions. The clean WikiHow dataset is split into 142,063 / 18,909 / 42,624 as train / dev / test sets. In order to retrieve the candidate answer pool for all the questions, we write a crawler to collect the relevant questions for each question from the *WikiHow* website. The answers of the relevant questions posted on *WikiHow* are labeled as negative answers for the given question. Finally, we obtain 1,188,189 question-answer pairs with corresponding answer summaries and matching labels as the WikiHowQA dataset. In accordance with the clean WikiHow dataset, we split the WikiHowQA dataset into 904,460 / 72,474 / 211,255 as train / dev / test sets, which implies that there is no overlapping of samples among the three split sets. The statistics of the WikiHowQA⁴ dataset are shown in Table 1.

In addition, we evaluate the proposed method on a resource-poor CQA dataset, StackExchange (Rücklé, Moosavi, and Gurevych 2019), which lacks of reference answer summaries. The statistics of the StackExchange dataset are presented in Table 2, which is a real-life CQA dataset containing data with long answers from different domains, including travel, cooking, academia, apple, and aviation. We adopt WikiHowQA as the source dataset for transfer learning due to its high quality and large quantity, while StackExchange are used as the target dataset.

Implementation Details

We train all the implemented models with pre-trained GloVe embeddings⁵ of 100 dimensions as word embeddings and set the vocabulary size to 50k for both source and target text. During training and testing procedure, we truncate the article to 400 words and restrict the length of generated summaries within 100 words. We apply early stopping based on the answer selection evaluation result on the validation set. We train our model and implement answer selection models for 5 epochs, while we implement summarization models for 20 epochs for fair comparisons, since the answers may repetitively occur in the candidates for different questions in the WikiHowQA dataset.

In our model, we train with a learning rate of 0.15 and an initial accumulator value of 0.1. The dropout rate is set to 0.5. The hidden unit sizes of the BiLSTM encoder and the LSTM decoder are all set to 150. We train our models with the batch size of 32. All other parameters are randomly initialized from $[-0.05, 0.05]$. $\lambda_1, \lambda_2, \lambda_3$ are all set to 1.

⁴<https://github.com/dengyang17/wikihowqa>

⁵<http://nlp.stanford.edu/data/glove.6B.zip>

Models	MAP	MRR
RANDOM GUESS	0.4088	0.4319
BM25	0.4212	0.4377
SIAMESE BiLSTM	0.4604	0.4734
ATT-BiLSTM	0.4573	0.4721
AP-BiLSTM	0.4896	0.5058
CA	<u>0.5022</u>	<u>0.5214</u>
COALA	0.5003	0.5196
GOLD + AP-BiLSTM	0.5261	0.5377
PGN + AP-BiLSTM	0.4992	0.5078
QPGN + AP-BiLSTM	0.5237	0.5343
QPGN + CA	0.5246	0.5373
QPGN + COALA	0.5197	0.5302
Joint Learning (ASAS)	0.5522	0.5686
w/o two-way attention	0.5208	0.5311
w/o pointer network	0.5341	0.5483

Table 3: Evaluation on Answer Selection

Experimental Result

Answer Selection Result

We first compare the proposed method with several state-of-the-art methods on the answer selection task, including Siamese BiLSTM (Mueller and Thyagarajan 2016), Att-BiLSTM (Tan et al. 2016), AP-LSTM (dos Santos et al. 2016), CA (Compare-Aggregate) (Wang and Jiang 2017) and COALA (Rücklé, Moosavi, and Gurevych 2019). Besides, we perform several **Two-Stage** methods, which first summarize the original answers and then conduct answer selection. To validate the effectiveness of different components of ASAS, we also conduct ablation tests. MAP and MRR are adopted as evaluation metrics.

Answer selection results on WikiHowQA are summarized in Table 3. We show that the joint learning model (ASAS) achieves state-of-the-art performance. There are several notable observations in the results. (i) BM25 model and even the basic deep learning model slightly improve the performance compared to random guessing, which signifies that the testing set is indeed a difficult one. (ii) The Compare-Aggregate methods (including CA and COALA) and AP-BiLSTM, which have been proven to be relatively effective in long-sentence answer selection (Rücklé, Moosavi, and Gurevych 2019; dos Santos et al. 2016), outperforms other strong baseline methods. (iii) Although **Two-Stage** methods actually improve the final answer selection result, it is time-consuming and inconvenient to train two separate models. In specific, using gold summary (GOLD) achieves the best performance, and Question-driven PGN (QPGN) performs better than original PGN. With the same summarization method, different answer selection models achieve similar results. (iv) Finally, the proposed joint learning model (ASAS) decently and substantially enhances the performance, which not only achieves the SOTA result, but also is easily trained by end-to-end fashion. By doing so, we precisely pick out the correct answers from candidate answers with long sentences, and meanwhile generate abstrac-

Models	ROUGE 1	ROUGE 2	ROUGE L
LEAD3	24.66	5.56	22.67
TEXTRANK	26.42	7.12	23.79
NEURALSUM	<u>27.01</u>	6.78	25.10
NEUSUM	26.78	6.88	25.14
SEQ2SEQ W/ ATTENTION	20.31	5.53	19.75
PGN W/ COVERAGE	26.83	<u>7.54</u>	<u>25.20</u>
SD ₂	26.65	6.92	24.77
BIASBLSTM	24.74	6.02	22.75
Question-driven PGN	27.32	7.98	25.46
Joint Learning (ASAS)	27.78	8.16	25.86

Table 4: Evaluation on Text Summarization

Method	Info	Conc	Read	Corr
NEURALSUM	3.60	2.70	3.22	3.24
PGN W/ COVERAGE	2.90	3.51	3.09	3.04
ASAS	3.67	3.88	3.59	3.71

Table 5: Human Evaluation Results

tive summaries for the convenience of community users. (v) The ablation study shows both the two-way attention mechanism and the pointer network contribute to the final result. The two-way attention mechanism enhances the interaction between questions and decoded answer summaries, while the pointer network aids in generating a better summary.

Answer Summary Generation Result

To evaluate the generated answer summary, we also compare the proposed method with the following state-of-the-art baseline methods on text summarization subtask, including four extractive methods (Lead3, TextRank (Mihalcea and Tarau 2004), NeuralSum (Cheng and Lapata 2016), NeuSum (Zhou et al. 2018)), two abstractive methods (Seq2Seq (Nallapati et al. 2016), PGN (See, Liu, and Manning 2017)) and two query-based methods (SD₂ (Nema et al. 2017), biASBLSTM (Singh et al. 2018)). ROUGE F1 scores are used to evaluate the summarization methods.

Text summarization results on WikiHowQA are summarized in Table 4. The experimental results show that the question-driven PGN outperforms all the state-of-the-art methods of both extractive and abstractive summarization, which demonstrates the effectiveness of incorporating question information to generate summaries for answers. The question information directly involves in the calculation of the generation probability to determine the next word whether generated from the vocabulary or copied from the source text. In addition, jointly learning with answer selection, ASAS further improves the result with a noticeable margin. The correlation information between question-answer pairs also aids in attending important words in the original answer, which are related to the question. These results show that ASAS can effectively generate high-quality summaries for the selected answers.

Analysis of The Length of Answers

In order to validate the effectiveness of the proposed method on long-sentence answer selection, we split the test set in

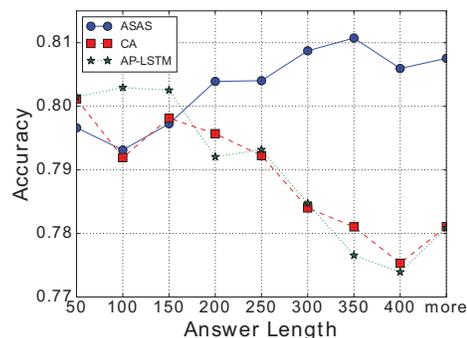


Figure 2: Model Accuracy in terms of Answer Length

QUESTION: How to <u>Donate Used Books to Charity?</u>
ANSWER: Many prisons offer books to prisoners ... <u>Donate Textbooks</u> is one organization that works with donated textbooks. (206 words)
REFERENCE SUMMARY: Donate books to prisoners. Donate books to soldiers. Donate books to underprivileged schools. Donate books internationally. Donate textbooks separately .
NEURALSUM: Do some research to find organizations that deliver books to prisons. Many prisons offer books to prisoners, which are received as donations. Consider books for adults with no violent or dark themes. Also, softcover books are preferred as many prisons do not allow hardcover books. Consider donating to books to prisoners. Books through bars also donates to prisons. Choose books that are entertaining and will provide some relaxation. Books for soldiers takes requests from soldiers and individuals are able to send those books to soldiers.
PGN W/ COVERAGE: <u>Find organizations</u> that deliver books. Choose books that are entertaining and entertaining. Research international donation organizations. <u>Find a organization</u> that works for accounting 101. <u>Find a organization</u> that works with textbook donations.
ASAS: Find organizations that work with prisons. Research international book organizations. <u>Donate books</u> to underprivileged schools. <u>Donate textbooks</u> for international book organizations. <u>Donate textbooks</u> to underprivileged. <u>Donate books</u> to prisons.

Figure 3: Case Study. ASAS generates the answer summary highly related to the question (Underlined), while PGN may misunderstand the core idea of the answer (Wavy-lined).

terms of the length of the answer. As shown in Fig. 2, we compare ASAS with two baseline methods, AP-LSTM and Compare-Aggregate Model (CA), by measuring the accuracy, which is the ratio of correctly selected answers. We observe that ASAS performs better especially for long answers. For answers that are shorter than 100 words, CA and AP-LSTM is slightly better than ASAS, which indicates that the summary may have lost some information for short answers. However, the performance of these two methods goes down with the increase in the answer length, while ASAS maintains a great stability.

Human Evaluation on Summarization

We conduct human evaluation on a sample of test set to evaluate the generated answer summaries from four aspects: (1) Informativity: how well does the summary capture the key information from the original answer? (2) Conciseness: how concise the summary is? (3) Readability: how fluent and coherent the summary is? (4) Correlatedness: how correlated the summary and the given question are? We randomly sample 50 answers and generate their summaries by three methods, including NeuralSum, PGN w/ coverage and the proposed ASAS. Three data annotators are asked to score each generated summary with 1 to 5 (higher the better).

Models	Finetune	Travel	Cooking	Academia	Apple	Aviation
BM25	-	38.1	30.9	29.2	21.8	37.0
BiLSTM	-	45.3	35.2	31.5	27.2	37.3
ATT-BiLSTM	-	43.0	36.2	31.2	24.7	33.9
AP-BiLSTM	-	38.8	32.2	27.3	22.9	34.5
CA	-	46.5	39.4	36.1	29.2	46.5
COALA	-	<u>53.8</u>	47.3	<u>42.2</u>	32.0	48.4
AP-BiLSTM	No	39.7	34.4	30.6	25.7	34.8
CA	No	33.4	28.1	21.4	21.2	31.5
COALA	No	35.6	32.2	24.5	22.8	37.2
AP-BiLSTM	Yes	44.9	38.1	36.7	29.1	46.3
CA	Yes	46.2	39.9	36.6	29.5	45.2
COALA	Yes	52.7	<u>49.2</u>	41.5	<u>32.4</u>	<u>49.9</u>
ASAS	-	54.8	48.1	42.8	32.6	50.1
ASAS	No	52.3	45.8	39.9	30.9	48.2
ASAS	Yes	56.5	52.8	44.4	35.1	52.9

Table 6: Evaluation on Resource-poor Answer Selection

Table 5 shows the human evaluation results. The results show that ASAS consistently outperforms other methods in all aspects. Noticeably, the proposed method learns well to generate answer summaries that are highly related to the given questions so there is a substantial margin on Correlatedness. In order to intuitively observe the advantage of the proposed method, we randomly choose one example to show the answer summary generation results. As shown in the Fig. 3, the extractive method (e.g., NeuralSum) selects important sentences from the original answer to form the answer summary, which still contains many insignificant or redundant information. The abstractive method (e.g., PGN) generates the answer summary from the vocabulary and the original answer, which may miss some key words and essential information. Upon these defects, the proposed joint learning method (ASAS) takes into account the information provided by the question to capture the core idea of the original answer and generate a precise summary. More importantly, unlike other methods, answer summaries are generated at the same time that the answers are selected.

Resource-poor CQA Results

To evaluate the transferring ability and applicability of the proposed method, we conduct experiments on the resource-poor CQA task with transfer learning. We also conduct several ablations that use no pre-training or no fine-tuning, including (i) *Finetune/-* is the baseline without pre-training, (ii) *Finetune/No* is trained with the training set of source data without fine-tuning on the target training data, (iii) *Finetune/Yes* is to first pre-train a model on the source data, and then use the learned parameters to initialize the model parameters for only fine-tuning the answer selection part on the target data. Following previous studies (Rücklé, Moosavi, and Gurevych 2019), we adopt the ratio of correctly selected answers as the evaluation metrics. Note that we use an unsupervised summarization method, TextRank (Mihalcea and Tarau 2004), to generate reference summaries roughly for *Finetune/-* settings with ASAS, since there is no reference summary in the original StackExchange dataset.

The experimental results show that even with the coarse reference summaries, ASAS (*Finetune/-*) achieves the best performance in 4 out of 5 domains, which demonstrates the applicability of the proposed joint learning framework. Un-

QUESTION (ACADEMIA): Etiquette for an initial meeting with a prospective advisor: What is expected of the student?
ANSWER: For an initial meeting, I would take an general idea of the kinds of things you 'd like to do in a particular advisor 's group , ... (250 words)
GENERATED SUMMARY: Consider the kinds of things you 'd like to do in advance. Make a list of things that will be funded with. Consider the idea of an advisor. Have a research group.
QUESTION (COOKING): 3:1 ratio of oil: vinegar, vinaigrette: How can I use this in soups?
ANSWER: Ruhlman's ratios are very useful for baking, where the ingredients play an important role in the physical structure of the end ... (315 words)
GENERATED SUMMARY: Consider the ingredients. Add fat and vinegar. Add oil to the soup. Add a squirt to the soup. Add a few drops of vinegar to the soup.
QUESTION (TRAVEL): Long layover in London, time to visit Olympics?
ANSWER: There are actually a number of different questions there, and while I can't tackle all of them I can take a stab at a few ... (302 words)
GENERATED SUMMARY: Consider the Olympic lanes. Get a travel card. Get a security check. Pay attention to the weather. Look at the train.

Figure 4: Generated Summaries for Resource-poor CQA

der the zero-shot setting, ASAS (*Finetune/No*) also achieves competitive results as those strong baseline methods, which shows the strong transferring ability of the proposed method and the value of the large-scale source dataset, WikiHowQA. Fine-tuning the answer selection part further outperforms all the baselines by about 4%. This result indicates that there are actually some gaps between different CQA datasets and the fine-tune strategy effectively overcomes these domain differences. Compared with ASAS and AP-BiLSTM, CA and COALA hardly benefit from pre-training due to their reliance on unsupervised embedding matching features.

In addition, Fig. 4 presents examples of answer summary generation results from target datasets. For those resource-poor CQA tasks without reference answer summaries, ASAS can not only achieve state-of-the-art results on answer selection, but also automatically generate decent and concise summaries via a simple transfer learning strategy with a resource-rich dataset.

Conclusion

We study the joint learning of answer selection and answer summary generation in CQA. We propose a novel model to employ the question information to improve the summarization result, and meanwhile leverage the summaries to reduce noise in answers for a better performance on long-sentence answer selection. In order to evaluate the answer generation task in CQA, we construct a new large-scale CQA dataset, WikiHowQA, which contains both labels for answer selection task and reference summaries for text summarization task. The experimental results show that the proposed joint learning method outperforms the state-of-the-art methods on both answer selection and summarization tasks, and processes robust applicability and transferring ability for resource-poor CQA tasks.

References

- Cheng, J., and Lapata, M. 2016. Neural summarization by extracting sentences and words. In *ACL*.
- Choi, E.; Hewlett, D.; Uszkoreit, J.; Polosukhin, I.; Lacoste, A.; and Berant, J. 2017. Coarse-to-fine question answering for long documents. In *ACL*, 209–220.

- Cohen, D.; Yang, L.; and Croft, W. B. 2018. Wikipassageqa: A benchmark collection for research on non-factoid answer passage retrieval. In *SIGIR*, 1165–1168.
- Deng, Y.; Shen, Y.; Yang, M.; Li, Y.; Du, N.; Fan, W.; and Lei, K. 2018. Knowledge as A bridge: Improving cross-domain answer selection with external knowledge. In *COLING*, 3295–3305.
- Deng, Y.; Xie, Y.; Li, Y.; Yang, M.; Du, N.; Fan, W.; Lei, K.; and Shen, Y. 2019. Multi-task learning with multi-view attention for answer selection and knowledge base question answering. In *AAAI*, 6318–6325.
- dos Santos, C. N.; Tan, M.; Xiang, B.; and Zhou, B. 2016. Attentive pooling networks. *CoRR*, abs/1602.03609.
- Joty, S. R.; Màrquez, L.; and Nakov, P. 2018. Joint multitask learning for community question answering using task-specific embeddings. In *EMNLP*, 4196–4207.
- Koupaee, M., and Wang, W. Y. 2018. Wikihow: A large scale text summarization dataset. *CoRR* abs/1810.09305.
- Li, Y.; Du, N.; Liu, C.; Xie, Y.; Fan, W.; Li, Q.; Gao, J.; and Sun, H. 2017. Reliable medical diagnosis from crowdsourcing: Discover trustworthy answers from non-experts. In *WSDM*, 253–261.
- Mihalcea, R., and Tarau, P. 2004. Textrank: Bringing order into text. In *EMNLP*, 404–411.
- Moschitti, A.; Bonadiman, D.; and Uva, A. 2017. Effective shared representations with multitask learning for community question answering. In *EACL*, 726–732.
- Mueller, J., and Thyagarajan, A. 2016. Siamese recurrent architectures for learning sentence similarity. In *AAAI*.
- Nakov, P.; Hoogeveen, D.; Màrquez, L.; Moschitti, A.; Mubarak, H.; Baldwin, T.; and Verspoor, K. 2017. Semeval-2017 task 3: Community question answering. In *SemEval@ACL*, 27–48.
- Nallapati, R.; Zhou, B.; dos Santos, C. N.; Gülçehre, Ç.; and Xiang, B. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. In *CoNLL*.
- Nallapati, R.; Zhai, F.; and Zhou, B. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *AAAI*, 3075–3081.
- Nema, P.; Khapra, M. M.; Laha, A.; and Ravindran, B. 2017. Diversity driven attention model for query-based abstractive summarization. In *ACL*, 1063–1072.
- Nishida, K.; Saito, I.; Nishida, K.; Shinoda, K.; Otsuka, A.; Asano, H.; and Tomita, J. 2019. Multi-style generative reading comprehension. In *ACL*, 2273–2284.
- Rücklé, A.; Moosavi, N. S.; and Gurevych, I. 2019. COALA: A neural coverage-based approach for long answer selection with small data. In *AAAI*, 6932–6939.
- Saha, A.; Aralikkatte, R.; Khapra, M. M.; and Sankaranarayanan, K. 2018. Duorc: Towards complex language understanding with paraphrased reading comprehension. In *ACL*, 1683–1693.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get to the point: Summarization with pointer-generator networks. In *ACL*, 1073–1083.
- Severyn, A., and Moschitti, A. 2015. Learning to rank short text pairs with convolutional deep neural networks. In *SIGIR*, 373–382.
- Shen, Y.; Deng, Y.; Yang, M.; Li, Y.; Du, N.; Fan, W.; and Lei, K. 2018. Knowledge-aware attentive neural network for ranking question answer pairs. In *SIGIR*, 901–904.
- Singh, M.; Mishra, A.; Oualil, Y.; Berberich, K.; and Klakow, D. 2018. Long-span language models for query-focused unsupervised extractive text summarization. In *ECIR*, 657–664.
- Song, H.; Ren, Z.; Liang, S.; Li, P.; Ma, J.; and de Rijke, M. 2017. Summarizing answers in non-factoid community question-answering. In *WSDM*, 405–414.
- Tan, M.; Santos, C. D.; Xiang, B.; and Zhou, B. 2016. Improved representation learning for question answer matching. In *ACL*, 464–473.
- Tomasoni, M., and Huang, M. 2010. Metadata-aware measures for answer summarization in community question answering. In *ACL*, 760–769.
- Wang, S., and Jiang, J. 2017. A compare-aggregate model for matching text sequences. In *ICLR*.
- Wang, M., and Manning, C. D. 2010. Probabilistic tree-edit models with structured latent variables for textual entailment and question answering. In *COLING*, 1164–1172.
- Wang, D., and Nyberg, E. 2015. A long short-term memory model for answer sentence selection in question answering. In *ACL*, 707–712.
- Wang, S.; Yu, M.; Guo, X.; Wang, Z.; Klinger, T.; Zhang, W.; Chang, S.; Tesauero, G.; Zhou, B.; and Jiang, J. 2018a. R³: Reinforced ranker-reader for open-domain question answering. In *AAAI*, 5981–5988.
- Wang, Z.; Liu, J.; Xiao, X.; Lyu, Y.; and Wu, T. 2018b. Joint training of candidate extraction and answer selection for reading comprehension. In *ACL*, 1715–1724.
- Wang, K.; Ming, Z.; and Chua, T. 2009. A syntactic tree matching approach to finding similar questions in community-based qa services. In *SIGIR*, 187–194.
- Wen, J.; Ma, J.; Feng, Y.; and Zhong, M. 2018. Hybrid attentive answer selection in CQA with deep users modelling. In *AAAI*, 2556–2563.
- Wu, W.; Sun, X.; and Wang, H. 2018. Question condensing networks for answer selection in community question answering. In *ACL*, 1746–1755.
- Yang, M.; Chen, L.; Chen, X.; Wu, Q.; Zhou, W.; and Shen, Y. 2019. Knowledge-enhanced hierarchical attention for community question answering with multi-task and adaptive learning. In *IJCAI*, 5349–5355.
- Yang, Y.; Yih, W. T.; and Meek, C. 2015. Wikiqa: A challenge dataset for open-domain question answering. In *EMNLP*, 2013–2018.
- Yoon, S.; Shin, J.; and Jung, K. 2018. Learning to rank question-answer pairs using hierarchical recurrent encoder with latent topic clustering. In *NAACL-HLT*, 1575–1584.
- Yu, J.; Qiu, M.; Jiang, J.; Huang, J.; Song, S.; Chu, W.; and Chen, H. 2018. Modelling domain relationships for transfer learning on retrieval-based question answering systems in e-commerce. In *WSDM*, 682–690.
- Zhou, Q.; Yang, N.; Wei, F.; Huang, S.; Zhou, M.; and Zhao, T. 2018. Neural document summarization by jointly learning to score and select sentences. In *ACL*, 654–663.
- Zhou, L.; Lin, C.; and Hovy, E. H. 2006. Summarizing answers for complicated questions. In *LREC*, 737–740.