

KNOWLEDGE DISTILLATION FOR IMPROVED ACCURACY IN SPOKEN QUESTION ANSWERING

Chenyu You^{1†} Nuo Chen^{2†} Yuexian Zou^{2,3,*}

¹Department of Electrical Engineering, Yale University, CT, USA

²ADSPLAB, School of ECE, Peking University, Shenzhen, China

³Peng Cheng Laboratory, Shenzhen, China

chenyu.you@yale.edu, {nuochen, zouyx}@pku.edu.cn

ABSTRACT

Spoken question answering (SQA) is a challenging task that requires the machine to fully understand the complex spoken documents. Automatic speech recognition (ASR) plays a significant role in the development of QA systems. However, the recent work shows that ASR systems generate highly noisy transcripts, which critically limit the capability of machine comprehension on the SQA task. To address the issue, we present a novel distillation framework. Specifically, we devise a training strategy to perform knowledge distillation (KD) from spoken documents and written counterparts. Our work aims at distilling rich knowledge from the language model to improve the performance of the student model by reducing the misalignment between automatic and manual transcripts. Experiments demonstrate that our approach outperforms several state-of-the-art language models on the Spoken-SQuAD dataset.

Index Terms— knowledge distillation, spoken question answering, question answering

1. INTRODUCTION

Text-based question answering (QA) [1, 2, 3] is an important task in natural language processing (NLP), which requires the machine to find the relevant textual answers in a context given a natural language question. In recent years, deep neural networks have achieved substantial progress in various QA applications. However, the voice interfaces for QA systems are less investigated. The main reason is the lack of large datasets with abundant annotations. Spoken-SQuAD [4] is one of the few available datasets for spoken question answering (SQA) tasks, which leverages the text documents to

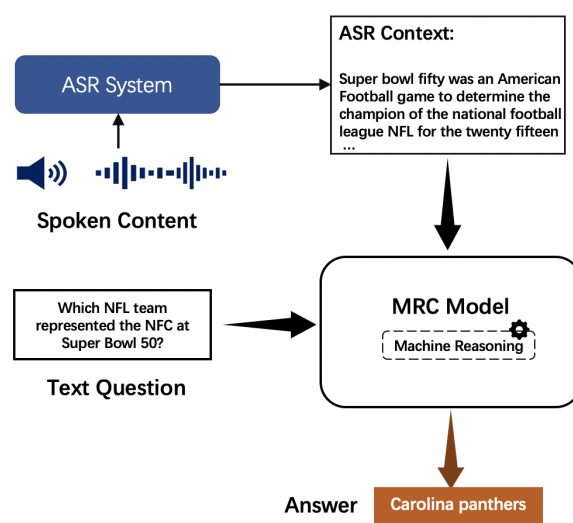


Fig. 1. Flow diagram of SQA systems on Spoken-SQuAD dataset.

generate corresponding spoken documents via Google Text-to-Speech system. Meanwhile, different from traditional text-based QA tasks, spoken question answering (SQA) includes audio signal processing, passage comprehension, and contextual understanding. Previous SQA approaches first transfer spoken content into text transcripts via ASR systems, then deploy some effective methods such as similarity matching [5], information retrieval [6] to predict answers given the ASR transcripts. Previous work has shown that ASR errors (e.g., Barcelona” to “bars alone”) make the transcribed content much more difficult for the QA systems to understand and make reasonable predictions.

The combination of ASR errors and loose syntax are the major impediments when handling spoken documents. To mitigate the impact of transcription errors, several methods [7, 8, 9, 10, 11, 12] have been proposed. Some recent methods [7, 8] adopt sub-word unit features strategy to enhance the word embedding with sub-word embedding to reduce unexpected ASR errors. Since ASR errors can be re-

[†] Indicates equal contribution.

This paper is partially supported by Shenzhen municipal research funding project “Research on key technologies of Spoken Language Understanding via Representation Learning and Knowledge Assisting”. Special acknowledgements are given to AOTO-PKUSZ Joint Lab for its support.

* Corresponding Author: zouyx@pku.edu.cn

garded as substitutions of word sequences for other ones (e.g., "feature" to "feather"), some sub-word sequences can still be correctly recognized when the corresponding transcribed content is wrong. However, it is still unclear whether we can explicitly learn relations in the discourse for the downstream tasks such as building QA systems. More recently, Lee et al. [10] proposes an adversarial domain adaptation method to reduce the gap between ASR hypotheses and the corresponding written content. Although the previous study has shown promising performance gains, it may suffer instability in training.

In this paper, we focus on the spoken question answering tasks. To resolve the challenges, we present a simple yet effective distillation approach, which fuses ASR-robust features to reduce the misalignment between ASR hypotheses and the reference transcripts. More concretely, we first take advantage of the dual nature property between speech utterances and text content by training the *teacher* model to learn two forms of the correspondences. We then distill this knowledge to the *student* model to mitigate the effect of unexpected transcript errors to boost the SQA performance.

2. METHOD

2.1. Task definition

Given a dataset $\mathcal{X} \in \{Q_i, D_i, A_i\}_i^N$, where Q_i is a question in text or spoken form, D_i is a document (passage) in spoken form and A_i is the answer in text form, respectively. Specifically, Q_i and A_i are both in the form of a single sentence, while D_i consists of multiple sentences. In this study, we use the Spoken-SQuAD dataset to validate the effectiveness of our approach on the extractive SQA task. Our goal is to find the text-based answer span A_i from the transcriptions of spoken document D_i given the question Q_i .

In this work, our QA system for spoken documents includes two sub-modules: ASR module and question answering module. The QA module predicts the answer given a written question and speech transcriptions. We present the flow diagram of SQA on Spoken-SQuAD dataset in Figure 1.

2.2. Knowledge Distillation

Hinton et al. [13] first introduced the idea of *knowledge distillation* (KD). In recent years, KD-based methods have achieved remarkable success in many computer vision and natural language processing tasks, such as natural language inference [14], sentiment classification [4], and object detection [15]. The key idea behind KD is to apply the prediction output of the *teacher* model as a 'soft label', and then utilize it as a supervision signal to guide the training of the *student* model. We set the model trained on the text transcripts as the *teacher* model $F_T(\cdot)$ and trained on the ASR hypotheses as the *student* model $F_S(\cdot)$, respectively.

Specifically, let x and y be the training input and ground-truth label from the dataset \mathcal{X} . Suppose the z_S and z_T are the output logits by $F_S(\cdot)$ and $F_T(\cdot)$, $F_S(\cdot)$ is trained to optimize the following objective:

$$L = \sum_{x \in \mathcal{X}} (\alpha \tau^2 \mathcal{KL}(p_\tau(z_S), p_\tau(z_T)) + (1 - \alpha) \mathcal{XE}(z_T, y)),$$

where $\mathcal{KL}(\cdot)$ and $\mathcal{XE}(\cdot)$ represent the Kullback-Leibler divergence and cross entropy, respectively. $p_\tau(\cdot)$ denotes the softmax function with temperature τ , and α is a balancing factor.

2.3. Models

To demonstrate the effectiveness and generality of our method, we choose two state-of-the-art QA models (BiDAF [16], QANet [17], and two pre-trained language models: BERT [18] and ALBERT [19]) as the baseline, which achieve superior performance on Text-SQuAD [20] dataset.

3. EXPERIMENTS

3.1. Dataset

Spoken-SQuAD (S-SQuAD) [7] is an English listening comprehension dataset, which contains 37,111 question pairs in the training set and 5,351 in the testing set. In Spoken-SQuAD, the documents are in spoken form, and the questions and answers are in the text form. The word error rate (WER) is 22.77% on the training set and 22.73% on the testing set. The reference written documents in Spoken-SQuAD are from Text-SQuAD [20], which is one of the most popular machine reading comprehension datasets. Text-SQuAD (T-SQuAD) includes 536 articles randomly selected from Wikipedia and the corresponding human-made 107,785 question-answer pairs.

In our experiments, we set our baselines trained on the Text-SQuAD as the *teacher* model and trained on the Spoken-SQuAD as the *student* model. Specially, we first train baselines on the Text-SQuAD training set and evaluate the performances of testing baselines on Text-SQuAD dev set and Spoken-SQuAD test set, respectively. Then, we train the baselines on the Spoken-SQuAD training set and evaluate the baselines on the Text-SQuAD dev set and Spoken-SQuAD test set, respectively. We report the quantitative results in Table 1.

3.2. Implementation Details

In our experiments, we use PTB tokenizer and spaCy tokenizer for BiDAF and QANet, respectively. For BiDAF and QANet, we adopt pre-trained word vectors, GloVe [21] word vectors to achieve the fixed word embedding of each word from both Text-SQuAD and Spoken-SQuAD. The batch size

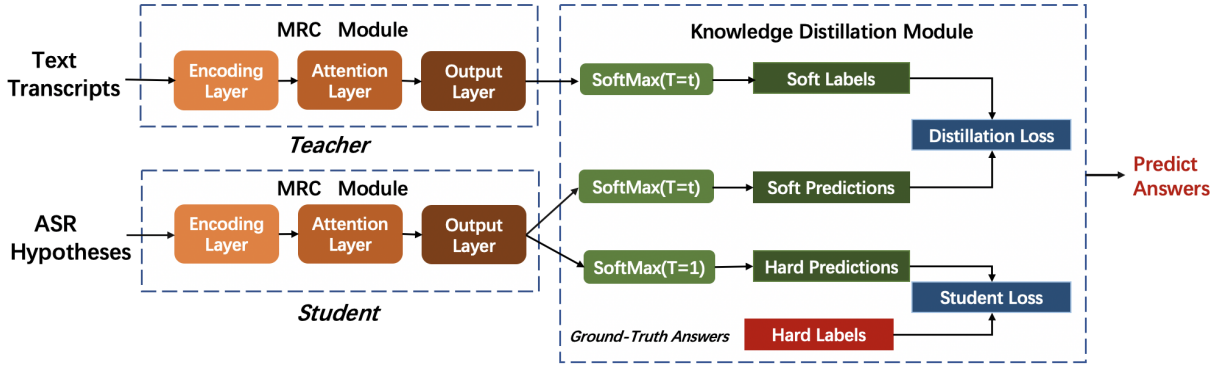


Fig. 2. An overview of our proposed framework.

Table 1. Comparison of four baselines (BiDAF, QANet, BERT, ALBERT). Note that, for brevity, we refer Text-SQuAD dev set and Spoken-SQuAD test set as T-SQuAD dev and S-SQuAD test, respectively.

Methods	T-SQuAD				S-SQuAD			
	T-SQuAD dev		S-SQuAD test		T-SQuAD dev		S-SQuAD test	
	EM	F1	EM	F1	EM	F1	EM	F1
BiDAF [16]	65.8	74.1	42.4	54.8	44.9	56.6	46.1	58.9
QANet [17]	70.9	80.5	49.5	63.2	51.1	65.5	54.5	66.1
BERT-large [18]	83.7	88.7	60.8	72.3	58.3	70.2	58.6	71.1
ALBERT-large [19]	84.4	89.6	61.6	73.8	59.1	71.9	59.4	72.2
Average	76.2	83.2	53.6	65.3	53.4	66.1	54.6	67.1

is set to be 10. We set the learning rate as $1e-3$. As for two pre-trained models (BERT and ALBERT), we use BPE tokenizer. We use a batch size of 2. The learning rate is set to be $2e-5$. We utilize the large version of BERT and ALBERT in our experiments, which composes of 24 transformer layers, and the hidden size of the word vector is 1024 in this study. We train four baselines (BiDAF, QANet, BERT, ALBERT) for 20, 15, 3, 3 epochs, respectively. To maintain the integrity of all evaluated model performance, we adopt standard implementations and hyper-parameters of four baselines for training. We set the balancing factor α as 0.9, and the temperature τ is 2. We use the F1 and Exact Match (EM) to evaluate SQA model performance.

4. RESULTS

To investigate the effects of ASR error on the SQA task, we compare four baselines. Table 1 reports the quantitative results. Table 1 shows that, in all baseline, the performance trained on ASR transcriptions is far lower than that of training on the text documents. This suggests that the word recognition errors affect significant content words, which lead the SQA model to make predictions incorrectly. Taking ALBERT-large as an example, when training on the T-SQuAD and testing on the T-SQuAD, it achieves 84.4%/89.6% on EM/F1 score. In contrast, when training on the S-SQuAD and testing on the S-SQuAD test set, ALBERT-large only obtains 59.4%/72.2% on EM/F1 score. This confirms that ASR errors severely degrade the performance of SQA models. Thus, it is necessary to explore how

to mitigate the impact of ASR errors for SQA systems.

To demonstrate the effectiveness of our proposed method, we compare four baseline methods *with* and *without* KD. We adopt the following setting - the *student* and *teacher* enjoy exactly the same network architecture. As shown in Table 2, it demonstrates that adopting the *knowledge distillation* strategy consistently boost the remarkable performance on all baselines. For BiDAF, our method achieves 46.5%/58.7% (vs.44.9%/56.6%) and 47.9%/60.0% (vs.46.1%/58.9%) in terms of EM/F1 score over the text documents and ASR transcriptions, respectively. For QANet, our method outperforms the baseline without distillation, achieving 52.7%/67.9% (vs.51.1%/65.5%) and 55.6%/67.7% (vs.54.5%/66.1%) in terms of EM/F1 score. As for two BERT-based models (BERT-large and ALBERT-large), our methods with KD consistently improve EM/F1 scores to 60.1%/72.2% (vs.58.3%/70.2%) and 58.6%/71.1% (vs.60.8%/73.8%); 60.8%/73.6% (vs.59.1%/71.9%) and 61.9%/74.7% (vs.59.4%/72.2%), respectively. These results confirm the importance of *knowledge distillation* in this setting.

To analyze the effect of the pre-trained language model on the SQA task, we first pick SDNet [2], which uses the pre-trained BERT-based and LSTM as its embedding encoder, and 283M parameters in total. For comparison, we establish the lightweight version of SDNet (17.7M parameters) by removing the pre-training language model (BERT-base) from SDNet. As shown in Table 3, we can observe that, on the Spoken-SQuAD, the original SDNet consistently outperforms the one without BERT. This suggests that the importance of the pre-trained language model on the SQA task. Meanwhile,




Fig. 3. Ablation studies of temperature τ on four baseline performances (BiDAF, QANet, BERT, ALBERT). Green and Orange denote the results of F1 score and EM score on Spoken-SQuAD test set, respectively.

Table 2. Comparison of our method. We set the model on text corpus as the *teacher* model, and the one on the ASR transcripts as the *student* model.

Methods	T-SQuAD dev		S-SQuAD test	
	EM	F1	EM	F1
BiDAF [16]	44.9	56.6	46.1	58.9
+ Sub-word [7]	45.8	57.7	47.1	59.4
+ Do-Adaptation [10]	45.9	57.4	47.0	59.1
+ KD	46.5	58.7	47.9	60.0
QANet [17]	51.1	65.5	54.5	66.1
+ Sub-word [7]	51.9	66.6	55.0	66.6
+ Do-Adaptation [10]	52.2	66.7	55.2	66.8
+ KD	52.7	67.9	55.6	67.7
BERT-base [18]	58.3	70.2	58.6	71.1
+ Sub-word [7]	59.3	71.1	59.7	72.4
+ Do-Adaptation [10]	59.5	71.2	59.8	72.6
+ KD	60.1	72.2	60.8	73.8
ALBERT-base [19]	59.1	71.9	59.4	72.2
+ Sub-word [7]	59.7	72.8	60.7	73.5
+ Do-Adaptation [10]	60.0	72.7	60.8	73.5
+ KD	60.8	73.6	61.9	74.7

to demonstrate the capability of knowledge distillation, we set the original SDNet trained on the Spoken-SQuAD as the massive *teacher* model, and SDNet without BERT trained on the Spoken-SQuAD as the lightweight *student* model. We find that performing *knowledge distillation* can indeed improve the performance of QA on speech transcripts. This suggests that distilling highly effective language representations learned by the large-scale *teacher* model to the small *student* model, which can improve the training process and the performance of the compact model.

5. QUANTITATIVE ANALYSIS

Effect Of Temperature τ . To investigate the impact of temperature τ , we evaluate the performances of four baselines with the standard choice of the temperature $\tau \in \{1, 2, 4, 6, 8, 10\}$. As shown in Figure 3, we find that four baselines consistently achieve their best performances in terms of F1 and EM scores with T setting to 2. All models are trained on Spoken-CoQA dataset, and validated on the Spoken-CoQA test set, respectively. It is noteworthy that we see the similar trends are emerging when testing baselines on

Table 3. Comparison the performance of SDNet [2] and SDNet [2] without (w/o) BERT with respect to manual transcriptions (Text-SQuAD) or ASR hypothesis (Spoken-SQuAD).

Model	Data Usage		Result	Parameters
	Training	Dev&Test	EM/F1	Millions
SDNet [2]	Manual	Manual	78.6/87.1	283
w/o BERT	Manual	Manual	72.5/80.3	17.7
SDNet	ASR	ASR	57.8/71.8	241.4
SDNet+ KD	ASR	ASR	59.2/73.6	241.4
w/o BERT	ASR	ASR	51.8/63.7	15.1
w/o BERT+ KD	ASR	ASR	52.2/64.5	15.1

Text-SQuAD dev set.

Generalizability for Knowledge Distillation. Table 2 shows that our proposed method significantly improves the network performance compared with the state-of-the-art baseline models on Spoken-SQuAD. As shown in Table 3, we examine whether the *knowledge distillation* can mitigate the impact of ASR errors to achieve better accuracy in both cases over the ASR transcripts. There is some evidence that using KD strategy behaves better than the same model without KD. As can be seen, as the *teacher* model is much larger than the *student* model, the *student* is able to mimic the *teacher*, suggesting *teachers* yield better *students* accuracy than a *student* trained alone. In other words, it suggests that performing *knowledge distillation* is a promising research direction in the spoken question answering task.

6. CONCLUSION

In this paper, we have presented a data distillation approach to mitigate the impact of ASR errors. In our *teacher-student* framework, we distill the superior knowledge from the model trained on textual documents to the one on ASR transcripts to achieve better performance by taking advantage of interactivity in spoken-written documents to reduce transcript errors and loose syntax. Experiments demonstrate the effectiveness of our proposed model, suggesting the generalizability for knowledge distillation. We believe that our data distillation technique may be a promising direction to advance different speech processing and natural language processing tasks.

7. REFERENCES

- [1] Hsin-Yuan Huang, Eunsol Choi, and Wen-tau Yih, “FlowQA: Grasping flow in history for conversational machine comprehension,” *arXiv preprint arXiv:1810.06683*, 2018.
- [2] Chenguang Zhu, Michael Zeng, and Xuedong Huang, “SDNet: Contextualized attention-based deep network for conversational question answering,” *arXiv preprint arXiv:1812.03593*, 2018.
- [3] Nuo Chen, Fenglin Liu, Chenyu You, Peilin Zhou, and Yuexian Zou, “Adaptive bi-directional attention: Exploring multi-granularity representations for machine reading comprehension,” *arXiv preprint arXiv:2012.10877*, 2020.
- [4] Zheng Li, Ying Wei, Yu Zhang, Xiang Zhang, and Xin Li, “Exploiting coarse-to-fine task transfer for aspect-level sentiment classification,” in *AAAI*, 2019, vol. 33, pp. 4253–4260.
- [5] Dina Obeid, Hugo Ramambason, and Cengiz Pehlevan, “Structured and deep similarity matching via structured and deep hebbian networks,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019, pp. 15403–15412.
- [6] Shao-Wei Fan-Jiang, Tien-Hong Lo, and Berlin Chen, “Spoken document retrieval leveraging bert-based modeling and query reformulation,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8144–8148.
- [7] Chia-Hsuan Li, Szu-Lin Wu, Chi-Liang Liu, and Hung-yi Lee, “Spoken SQuAD: A study of mitigating the impact of speech recognition errors on listening comprehension,” *arXiv preprint arXiv:1804.00320*, 2018.
- [8] Chia-Hsuan Lee, Shang-Ming Wang, Huan-Cheng Chang, and Hung-Yi Lee, “ODSQA: Open-domain spoken question answering dataset,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 949–956.
- [9] Chenyu You, Nuo Chen, and Yuexian Zou, “Contextualized attention-based knowledge transfer for spoken conversational question answering,” *arXiv preprint arXiv:2010.11066*, 2020.
- [10] Chia-Hsuan Lee, Yun-Nung Chen, and Hung-yi Lee, “Mitigating the impact of speech recognition errors on spoken question answering by adversarial domain adaptation,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2019, pp. 7300–7304, IEEE.
- [11] Dan Su and Pascale Fung, “Improving spoken question answering using contextualized word representation,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8004–8008.
- [12] Chenyu You, Nuo Chen, Fenglin Liu, Dongchao Yang, and Yuexian Zou, “Towards data distillation for end-to-end spoken conversational question answering,” *arXiv preprint arXiv:2010.08923*, 2020.
- [13] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [14] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu, “Tinybert: Distilling bert for natural language understanding,” *arXiv preprint arXiv:1909.10351*, 2019.
- [15] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker, “Learning efficient object detection models with knowledge distillation,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 742–751.
- [16] Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi, “Bidirectional attention flow for machine comprehension,” *arXiv preprint arXiv:1611.01603*, 2016.
- [17] Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le, “Qanet: Combining local convolution with global self-attention for reading comprehension,” *arXiv preprint arXiv:1804.09541*, 2018.
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [19] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut, “Albert: A lite bert for self-supervised learning of language representations,” *arXiv preprint arXiv:1909.11942*, 2019.
- [20] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang, “SQuAD: 100,000+ questions for machine comprehension of text,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016, pp. 2383–2392.
- [21] Jeffrey Pennington, Richard Socher, and Christopher D Manning, “Glove: Global vectors for word representation,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543.