

Robust Domain Adaptation for Machine Reading Comprehension

Liang Jiang^{*1}, Zhenyu Huang^{*2}, Jia Liu¹, Zujie Wen¹, Xi Peng^{2†}

¹Ant Group

²College of Computer Science, Sichuan University
{tianxuan.jl, jianiu.lj, zujie.wzj}@antgroup.com, {zyhuang.gm, pengx.gm}@gmail.com

Abstract

Most domain adaptation methods for machine reading comprehension (MRC) use a pre-trained question-answer (QA) construction model to generate pseudo QA pairs for MRC transfer. Such a process will inevitably introduce mismatched pairs (i.e., Noisy Correspondence) due to i) the unavailable QA pairs in target documents, and ii) the domain shift during applying the QA construction model to the target domain. Undoubtedly, the noisy correspondence will degenerate the performance of MRC, which however is neglected by existing works. To solve such an untouched problem, we propose to construct QA pairs by additionally using the dialogue related to the documents, as well as a new domain adaptation method for MRC. Specifically, we propose Robust Domain Adaptation for Machine Reading Comprehension (RMRC) method which consists of an answer extractor (AE), a question selector (QS), and an MRC model. Specifically, RMRC filters out the irrelevant answers by estimating the correlation to the document via the AE, and extracts the questions by fusing the candidate questions in multiple rounds of dialogue chats via the QS. With the extracted QA pairs, MRC is fine-tuned and provides the feedback to optimize the QS through a novel reinforced self-training method. Thanks to the optimization of the QS, our method will greatly alleviate the noisy correspondence problem caused by the domain shift. To the best of our knowledge, this could be the first study to reveal the influence of noisy correspondence in domain adaptation MRC models and show a feasible solution to achieve the robustness against the mismatched pairs. Extensive experiments on three datasets demonstrate the effectiveness of our method.

Introduction

In recent, a number of domain adaptation (DA) methods (Cao et al. 2020; Wang et al. 2019; Lewis, Denoyer, and Riedel 2019) for machine reading comprehension (MRC) have been proposed, which usually pre-train an MRC model in a high-resource domain and then transfer it to a low-resource domain. Specifically, most existing methods (Cao et al. 2020; Wang et al. 2019; Lewis, Denoyer, and Riedel 2019) consist of two steps. First, they construct some pseudo QA pairs by using a pre-trained QA construction model from

the available documents in the target domain. Then they fine-tune the pre-trained MRC model by using the constructed pairs.

Although these methods have achieved promising results, almost all of them ignore the mismatched QA pairs, i.e., the irrelevant QA pairs which are wrongly treated as positive. In the scenario of MRC, such a so-called noisy correspondence (NC) issue (Huang et al. 2021; Yang et al. 2021, 2022) is caused by the following reasons. First, the domain adaptation methods for MRC often construct pseudo QA pairs by using the documents in the target domain which does not contain natural questions. As a result, the generated questions will be probably irrelevant to the answers. Second, the QA construction model is pre-trained in the source domain and then directly applied to the target domain without fine-tuning. In consequence, such a domain shift issue will lead to noisy correspondence. A toy example of NC is shown in Fig. 1(a), and more real-world samples generated by the existing works refer to Fig. 4. Notably, the NC problem is remarkably different from the well-studied noisy labels. To be specific, noisy labels generally refer to the category-level annotation errors of a given data point, whereas here NC refers to the mismatched relationship between two data points. Undoubtedly, NC will degenerate the performance of the MRC model, which is however neglected so far as we know.

To sum up, domain adaptation methods for MRC will face the NC challenge caused by i) the unavailable QA pairs in the target documents, and ii) domain shift during applying the QA construction model to the target domain. In this paper, we propose to construct more credible QA pairs by additionally using document-associated dialogues to fill the vacancy of the natural question information in the document. In addition, our method will fine-tune the QA construction model in the target domain with the help of the MRC model feedback. For clarity, we summarize the major differences between our MRC method and the existing ones in Fig. 1(b) by taking the customer service as a showcase without loss of generality. As shown, in real-world applications, the customer service usually answers customers' questions by referring to the documents, forming an associated dialogue. Existing works only use the sole document for QA construction and do not further fine-tune the QA construction model in the target domain. In contrast, our work leverages both documents and the associated dialogues for QA con-

^{*}These authors contributed equally.

[†]Corresponding author.

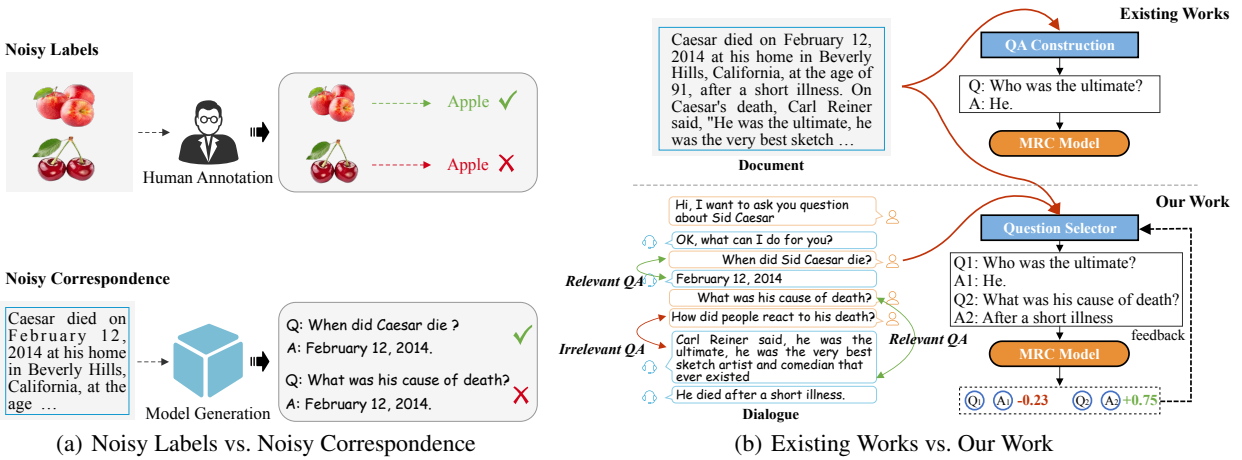


Figure 1: (a) Noisy Labels vs. Noisy Correspondence. The noisy labels refer to the errors in the category annotation of data samples caused by human annotation. The noisy correspondence here refers to the misalignment between two data points generated by the model itself. (b) An example of the document along with dialogue, and the difference between our method and the existing ones. The dialogue is the conversations between the questioner and the answerer (the customer and customer service here) about the document. In other words, the questioner raises a question about the document, and the answerer answers it by referring to the document. Hence, the dialogues have natural QA-form conversations that is helpful in QA pair construction.

struction. Dialogues are the conversation corpus between the questioners and answerers, which naturally preserves QA-similar chats and thus are more credible for QA construction. Moreover, our method will use the feedback of the MRC model on constructed QA pairs to optimize the QA construction model, thus alleviating the domain shift issue and improving the QA quality.

In practice, however, difficulties arise when an attempt is made to apply the above approaches. Specifically, although dialogues are more credible than documents for the QA construction, they still contain a huge number of irrelevant and discontinuous conversations. In other words, the above QA construction method could partially alleviate but still cannot solve the NC problem. As shown in Fig. 1(b), the irrelevant conversation is about the irrelevant conversations to the document, e.g., there are some greetings or other chats which are irrelevant w.r.t the document. The discontinuous conversation is that the question and answer are not exactly aligned in a single round of chat due to the complexity in interaction, e.g., the customer may raise a new question before receiving the answer for the last ones. Besides the challenges rooted in the dialogue data, another difficulty is how to fine-tune the non-differentiable QA construction model in the target domain to alleviate the domain shift issue. In brief, the existing works often generate QA pairs by resorting to the discrete $\arg \max$ sampling operator which hinders the optimization of the QA construction model.

To overcome the above challenges in data quality and model optimization, we propose a novel domain adaptation method for MRC, dubbed Robust Domain Adaptation for Machine Reading Comprehension (RMRC) which consists of an answer extractor (AE), a question selector (QS), and an MRC model. In brief, RMRC leverages both documents and the associated dialogue for MRC transfer by i) con-

structing QA pairs via the AE and the QS, and ii) alternative training the MRC model and the QS. In the first stage, for a given document, RMRC extracts the candidate answers from dialogues and filters out unrelated ones in terms of the estimated relevance using the pre-trained AE, thus alleviating the NC issue caused by irrelevant chats. After that, for the extracted answer, RMRC seeks the most related questions that appeared in multiple rounds of chats via the pre-trained QS, thus tackling the NC problem caused by discontinuous conversations. In the second stage, RMRC optimizes the MRC model and the QS in an alternative fashion, which will favor MRC transfer. In detail, the domain shift could be alleviated and accordingly the NC problem is tackled by optimizing the QS using the feedback of the MRC model. Note that, as the QS is non-differentiable and cannot be directly optimized by back-propagation, we propose a novel reinforced self-training optimization method by recasting the model evaluation on the constructed QA pairs as a training reward.

The main contributions and novelty of this paper could be summarized as below.

- This work is the first successful attempt to study the NC problem that is common but ignored in existing domain adaptation methods for MRC.
- To solve the NC problem, we propose to leverage both the document and associated dialogue for MRC model training. To the best of our knowledge, this could be the first method to leverage the dialogue in the domain adaptation for MRC.
- To implement robust domain adaptation method for MRC, RMRC consists of the AE, QS, and MRC model. Thanks to the reinforced self-training optimization method, the QS could be fine-tuned with the MRC

model feedback, thus further alleviating the influence of NC.

Related Work

In this section, we will briefly introduce some recent developments in machine reading comprehension, the domain adaptation for MRC, and noisy label learning.

Machine Reading Comprehension

Machine reading comprehension aims to read the documents and answer questions about a document. Thanks to the collected benchmark datasets like SQuAD (Rajpurkar et al. 2016), CoQA (Reddy, Chen, and Manning 2019), and QuAC (Choi et al. 2018), MRC has made great progress in recent years, and even surpasses the human level (Seo et al. 2016; Yu et al. 2018; Devlin et al. 2018; Zhang, Yang, and Zhao 2020; Jiang and Zhao 2018; Gao et al. 2019). (Seo et al. 2016) proposed BiDAF which leverages RNN and bi-directional attention mechanism between question and document to achieve promising performance in machine reading comprehension. QANet (Yu et al. 2018) used CNN rather than RNN to better capture local information in question and documents. BERT (Devlin et al. 2018) proposes to use a large amount of unsupervised corpus for pre-training and successfully improve performance on many downstream NLP tasks including MRC.

Domain Adaptation

Domain adaptation (DA) is a well-developed technique that aims to transfer knowledge in the presence of the domain gap, i.e. making a model pre-trained in the source domain generalizable to the target domain (Zhang et al. 2018, 2021; Ding et al. 2018; Hu et al. 2018; Kan, Shan, and Chen 2015; Li et al. 2013). The existing domain adaptation methods for MRC could be roughly grouped into two categories:

1) **Model generalization.** It aims to improve the generalization capability of the model trained on the source domain to the target domain (Li et al. 2019; Su et al. 2019; Baradaran and Amirkhani 2021). For example, (Su et al. 2019) propose to use a pre-trained model trained on multiple MRC datasets simultaneously to improve the generalization ability on the new domains. (Baradaran and Amirkhani 2021) propose to ensemble the models individually trained on different datasets to improve the generalization ability. However, this kind of methods do not leverage the available information on the target domain, thus obtaining less satisfactory performance.

2) **QA generation.** It often utilizes the target-domain documents to construct QA pairs to fine-tune the MRC models (Du, Shao, and Cardie 2017; Wang et al. 2019; Lewis, Denoyer, and Riedel 2019; Cao et al. 2020). For example, (Du, Shao, and Cardie 2017) propose to generate questions on the target domain using a Seq2seq question generator that trained on the source domain, and then fine-tune the MRC model with the pseudo QA pairs. (Cao et al. 2020) propose to add a self-training step to filter out low quality QA pairs using the estimated scores of each constructed pairs through the MRC model. Though these methods have made

great progress, they still suffer from the noisy correspondence problem as discussed in the introduction, thus leading to sub-optimal performance in real-world applications. Unlike these works, to address the NC problem, this paper proposes to construct more credible QA pairs by additionally using the dialogue, and optimize the MRC model and the question selector in an alternative fashion to alleviate the domain shift of the QS to the target domain.

Learning with Noisy Labels

To alleviate even eliminate the influence of noisy labels, many methods have been proposed to achieve robust classification results (Arazo et al. 2019; Song et al. 2020; Liu and Tao 2015; Xia et al. 2020; Bai et al. 2021; Luo, Han, and Gong 2021). Currently, the existing works often resort to sample selection for achieving noise robustness. To be specific, sample selection methods seek to select the clean samples from the noisy dataset for training. For example, (Arazo et al. 2019) proposes to select the samples with small loss as the clean samples. Moreover, to further enhance the clean sample selection capacity, Co-teaching methods (Han et al. 2018; Yu et al. 2019) leverage two individual trained networks to filter out noises in an alternative manner. In recent, PES (Bai et al. 2021) proposes a progressive early stopping strategy in the semi-supervised learning framework by treating the clean and noisy samples as the labeled and unlabeled data. In addition, some very recent works (Yang et al. 2021; Huang et al. 2021; Yang et al. 2022) study the paradigm of noisy correspondence like this paper. Although the works share some similarity with this work, they are significantly different in motivation, application, and method.

Method

In this section, we elaborate on the proposed Robust Domain Adaptation for Machine Reading Comprehension (RMRC).

Problem Definition

For a given document and a related question, MRC aims to find a text span from the document as the corresponding answer. To overcome the data scarcity issue in MRC, some domain adaptation methods are proposed to transfer an MRC model pre-trained in the source domain \mathcal{S} to the target domain \mathcal{T} . Different from existing domain adaptation methods for MRC, we utilize the documents along with the associated dialogues in \mathcal{T} to construct QA pairs for model fine-tuning. Formally, given a pre-trained MRC model M and B documents $\mathcal{D} = \{D_1, D_2, \dots, D_B\}$ in the target domain, where each document D_b is associated with a dialogue set generated by the questioners and answerers, we aim to transfer M from \mathcal{S} to \mathcal{T} . For each dialogue $C = \{C_1^{r_1}, C_2^{r_2}, \dots, C_{N_c}^{r_{N_c}} | r_i \in (a, q)\}$, it contains N_c chats, where r_i denotes the speaker of the i -th chat, a and q indicate the role of answerer and the questioner, respectively.

For clarity, we first provide an overview of the proposed RMRC and then introduce its components one by one. As shown in Fig. 2, RMRC consists of an answer extractor (AE), a question selector (QS), and an MRC model which

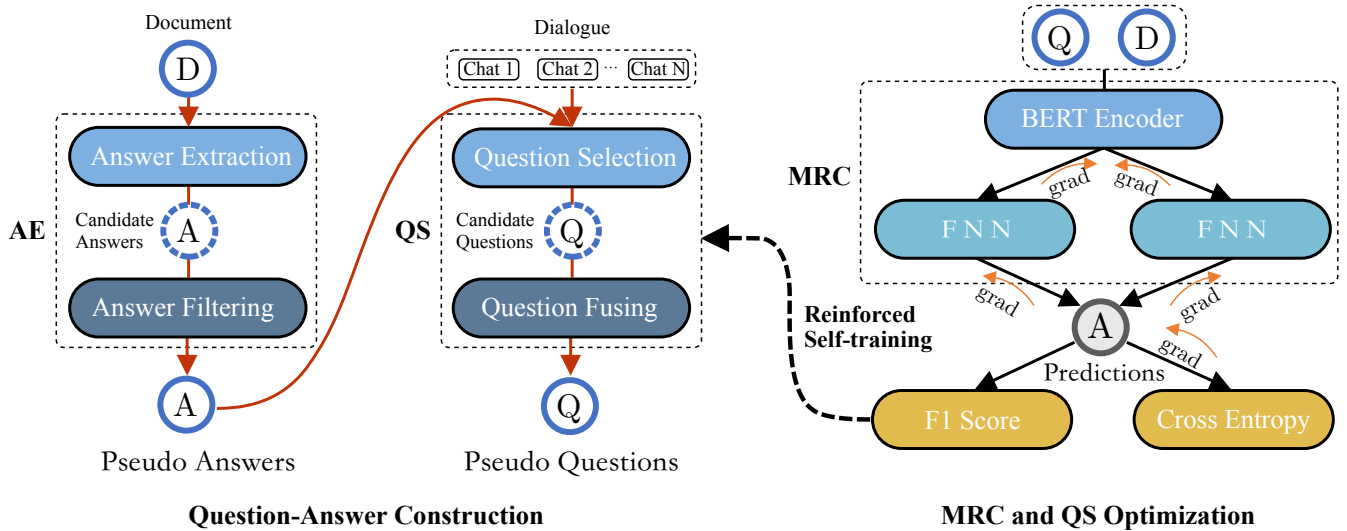


Figure 2: Overview of our method. RMRC consists of an answer extractor (AE), a question selector (AS), and an MRC model. In the first stage, AE and QS construct the QA pairs in the target domain by leveraging both the documents and dialogues. In the second stage, RMRC optimizes the MRC model and QS in an alternative fashion.

are applied to the following two stages: 1) **QA construction with AE and QS**. For each given document and the associated dialogues, the AE first splits the document into a set of candidate answers. Then, the AE outputs the pseudo answers by filtering out the irrelevant answers in terms of the relevance to the corresponding chat. After that, for each given pseudo answer and the related chat, QS selects multiple chats located before the answer-related chat in the dialogue as the candidate questions. Finally, the most related candidate questions are concatenated as pseudo questions. 2) **Alternative training of the MRC model and QS**. For a given pseudo question and document, RMRC uses the MRC model to obtain the corresponding answer. After that, the obtained answer is used to optimize the MRC model via a cross-entropy loss and the QS via our novel reinforce self-training optimizer.

Answer Extractor

The AE is designed to extract answers by finding the most similar text span from the document in terms of the answerer’s response contained in the dialogue. Formally, with the document set \mathcal{D} and associated dialogue \mathcal{C} , AE extracts the corresponding answer $\{\mathbf{a}_i\}_{i=1}^N \in \mathcal{D}$ by filtering out the unrelated answers. Specifically, we first split all documents in \mathcal{D} into n-gram tokens G via

$$G = \{G_1 \cup \dots \cup G_B\}$$

$$G_b = \{G_b^k\}_{k=1}^K, G_b^k = \text{n-gram}(D_b, k), \quad (1)$$

where n-gram is the operator to extract the candidate answers from \mathcal{D} , K denotes the maximum token number of the candidate answers, G_b^k denotes the k -gram token set from document G_b . In other words, we will extract 1-gram, 2-gram, \dots , K -gram tokens for answer extraction. With the candidate answers, a text matching function is developed to find the best matched answer $\mathbf{g}_i \in G$ to C_t^a which

denoted as \mathbf{a}_t , i.e.,

$$\mathbf{a}_t, D_t^* = \arg \max_{\mathbf{g}_i \in G_b, G_b \in G} S(C_t^a, \mathbf{g}_i), \quad (2)$$

where \mathbf{g}_i is the extracted n-gram token in G , (\mathbf{a}_t, D_t^*) are the best matched answer and the corresponding document to C_t^a . In the equation, we compute the cosine similarity $S(C_t^a, \mathbf{g}_i)$ between the extracted features of C_t^a and \mathbf{g}_i from a pre-trained BERT (Devlin et al. 2018). However in the real-world dialogues, the conversation contains not only the document-related answers but also some unrelated answers like greetings as shown in Fig. 1(b). Hence, to filter out these unrelated answers, we select the answer \mathbf{a}_t whose matching score is larger than a given threshold γ . Formally,

$$A = \{(\mathbf{a}_t, D_t^*, C_t^a) | \gamma \leq S(C_t^a, \mathbf{a}_t), \forall C_t^a \in \mathcal{C}\}. \quad (3)$$

Question Selector

For a given answer \mathbf{a}_t , we pass it and its associated document D_t^* along with the corresponding chat C_t^a through the QS to find the corresponding question \mathbf{q}_t from the dialogue. In detail, we first obtain the candidate questions by selecting multiple chats located before C_t^a in the dialogue, i.e.,

$$\tilde{Q} = \{C_{t'}^q | C_{t'}^q \in \mathcal{N}^\tau(C_t^a)\}, \quad (4)$$

where $\mathcal{N}^\tau(C_t^a)$ denotes τ closest chats located before C_t^a from questioner, τ is the maximal question selection range which is fixed to 16 throughout our experiments. Such a question selection strategy is based on the observation that the corresponding questions only exist before the answers and are more irrelevant as they are further away from the answer. For each $C_{t'}^q \in \tilde{Q}$, we then compute the relevance score $R(C_t^a, C_{t'}^q)$ between the answer C_t^a and the question $C_{t'}^q$ by

$$R(C_t^a, C_{t'}^q) = \text{Sigmoid}(h_{t'}^{qa})$$

$$h_{t'}^{qa} = \text{Encoder}_Q([C_t^a; C_{t'}^q]) \quad (5)$$

where Encoder_Q transforms $[C_t^a; C_{t'}^q]$ into a hidden vector for relevance prediction, $[\cdot]$ denotes the concatenation operator. Note that, the relevance score $R(C_t^a, C_{t'}^q)$ could be regarded as the conditional probability of $C_{t'}^q$ to be the corresponding question of the given answer C_t^a , i.e., $P(C_{t'}^q|C_t^a) \Leftrightarrow R(C_t^a, C_{t'}^q)$. To further alleviate the aforementioned discontinuous conversation issue in the real world (e.g., the questioner may raise a new question before the last one is answered.), we concatenate the most related questions in terms of the relevance score w.r.t. C_t^a , i.e.,

$$\mathbf{q}_t = \text{Concat}(\{C_{t'}^q|C_{t'}^q \in \mathcal{N}^\kappa(C_t^a)\}), \quad (6)$$

where $\mathcal{N}^\kappa(C_t^a)$ denotes κ -nearest neighbors of C_t^a in \tilde{Q} in terms of the relevance score. Accordingly, the probability of \mathbf{q}_t to be the corresponding question of C_t^a is formulated as,

$$P(\mathbf{q}_t|C_t^a) = \prod_{C_{t'}^q \in \mathcal{N}^\kappa(C_t^a)} P(C_{t'}^q|C_t^a). \quad (7)$$

MRC Model Training

With the constructed pseudo QA pairs and the corresponding document $\{(\mathbf{a}_i, \mathbf{q}_i, D_i)\}_{i=1}^N$, we fine-tune the pre-trained MRC model to improve its generalization to the target domain. In detail, we first embed the question and associated document into a hidden space with a BERT encoder, i.e.,

$$\mathbf{v}_i = \text{Encoder}_{\text{MRC}}([D_i; \mathbf{q}_i]). \quad (8)$$

Then, \mathbf{v}_i is utilized to predict the positions of the corresponding answer in the document. Specifically,

$$\begin{aligned} P_i^s &= \text{Softmax}(\text{FFN}(\mathbf{W}^s; \mathbf{v}_i)), \\ P_i^e &= \text{Softmax}(\text{FFN}(\mathbf{W}^e; \mathbf{v}_i)), \end{aligned} \quad (9)$$

where P_i^s and P_i^e denote the probability of each token being the start and end positions of the answer in D_i , respectively. $\text{FFN}(\mathbf{W}^s; \mathbf{v})$ and $\text{FFN}(\mathbf{W}^e; \mathbf{v})$ are two one-layer feed-forward networks with parameters \mathbf{W}^s and \mathbf{W}^e , respectively. Finally, we use the cross entropy between the predicted probability (P_i^s, P_i^e) and the ground truth as the training loss for our MRC model, i.e.,

$$\mathcal{L}_{\text{MRC}} = \frac{1}{N} \sum_i (\text{CE}(P_i^s, Y_i^s) + \text{CE}(P_i^e, Y_i^e)), \quad (10)$$

where Y_i^s and Y_i^e are two one-hot vectors which denote the start and end positions of \mathbf{a}_i in D_i .

Reinforced Self-training for QS

As the QS is pre-trained on the source domain, the noisy correspondence will be inevitably introduced when the QS is applied to the target domain for constructing QA pairs. To address this domain shift issue, we propose to optimize the QS in the target domain by recasting the model evaluation on the constructed pseudo QA pairs as a training reward. Specifically, for a given pair $(\mathbf{a}_i, \mathbf{q}_i, D_i)$, we first obtain the predicted answer by

$$\begin{aligned} \tilde{\mathbf{a}}_i &= [t_{\tilde{s}} : t_{\tilde{e}}], t \in D, \\ \tilde{s} &= \arg \max P_i^s, \\ \tilde{e} &= \arg \max P_i^e, \end{aligned} \quad (11)$$

where P_i^s and P_i^e are the model outputs defined in Eq. 9, $t_{\tilde{s}}$ and $t_{\tilde{e}}$ denote the start and end token of the predicted answer, respectively. By using the f1-score as the quality evaluation score of the constructed QA pairs, it is expected to generate QA pairs with high f1-score, i.e.,

$$\theta^* = \arg \max_{\theta} \mathbb{E} f1(\mathbf{a}, \tilde{\mathbf{a}}), \quad (12)$$

where θ denotes the parameters of QS. As shown in Eq. 6, the QA construction adopts the selection and concatenate operators via $\arg \max$ sampling technique which is discrete. To overcome this non-differentiable problem, we adopt policy gradient based reinforcement learning (REINFORCE) (Williams 1992) to approximate the gradients w.r.t. θ . Specifically, let $\mathcal{J}(\theta)$ denote the objective function (Eq. 12), its gradient can be approximated as below:

$$\begin{aligned} \nabla_{\theta} \mathcal{J}(\theta) &= \nabla_{\theta} \mathbb{E}_{\mathbf{q} \sim P_{\theta}(\mathbf{q})} f1(\mathbf{a}, \mathbf{f}(\mathbf{q})) \\ &= \sum_{i=1}^M \nabla_{\theta} P_{\theta}(\mathbf{q}_i) f1(\mathbf{a}_i, \mathbf{f}(\mathbf{q}_i)) \\ &= \sum_{i=1}^M P_{\theta}(\mathbf{q}_i) \nabla_{\theta} \log P_{\theta}(\mathbf{q}_i) f1(\mathbf{a}_i, \mathbf{f}(\mathbf{q}_i)) \\ &= \mathbb{E}_{\mathbf{q} \sim P_{\theta}(\mathbf{q})} [\nabla_{\theta} \log P_{\theta}(\mathbf{q}) f1(\mathbf{a}, \mathbf{f}(\mathbf{q}))] \\ &\approx \frac{1}{N} \sum_{i=1}^N \nabla_{\theta} \log P_{\theta}(\mathbf{q}_i) f1(\mathbf{a}_i, \mathbf{f}(\mathbf{q}_i)) \\ &= \frac{1}{N} \sum_{i=1}^N [\nabla_{\theta} \log \prod_{C_j^q \in \mathcal{N}_t^{\kappa}} P_{\theta}(C_j^q|C_t^a) f1(\mathbf{a}_i, \mathbf{f}(\mathbf{q}_i))] \\ &= \frac{1}{N} \sum_{i=1}^N [\nabla_{\theta} \sum_{C_j^q \in \mathcal{N}_t^{\kappa}} \log P_{\theta}(C_j^q|C_t^a) f1(\mathbf{a}_i, \mathbf{f}(\mathbf{q}_i))], \end{aligned} \quad (13)$$

where $\mathbf{f}(\mathbf{q})$ transforms the question \mathbf{q} into the predicted answer $\tilde{\mathbf{a}}$ via the MRC model, and \mathcal{N}_t^{κ} denotes $\mathcal{N}^{\kappa}(C_t^a)$ for simplicity, and $P_{\theta}(C_j^q|C_t^a)$ is conditional probability of C_j^q to be the corresponding question of C_t^a .

With the approximated gradients $\nabla_{\theta} \mathcal{J}(\theta)$, the loss function for QA could be rewritten by,

$$\mathcal{L}_Q = -\frac{1}{N} \sum_{i=1}^N \sum_{C_j^q \in \mathcal{N}_t^{\kappa}} \log R_{\theta}(C_t^a, C_j^q) [f1(\mathbf{a}_i, \mathbf{f}(\mathbf{q}_i)) - r_b], \quad (14)$$

where $R_{\theta}(C_t^a, C_j^q)$ denotes the relevance defined in Eq. 6. As only one sample per reward estimation is used, we use r_b as the baseline score for reducing variance like (Williams 1992).

Experiments

In this section, we evaluate the RMRC on three datasets by comparing it with three MRC domain adaptation methods.

Datasets

We pretrain the MRC model on the SQUAD dataset (Rajpurkar et al. 2016) and fine-tune it on three target datasets

Data	Size	BERT-S		UQACT		CAsE		AdaMRC		RMRC-fix		RMRC	
		EM	F1	EM	F1	EM	F1	EM	F1	EM	F1	EM	F1
d-QuAC	ALL	3.36	16.36	6.81	19.05	6.79	19.33	5.44	18.87	7.45	23.13	8.27	24.46
	5000	-	-	4.26	17.88	4.85	17.92	4.57	18.01	5.85	21.25	6.37	22.27
	1000	-	-	3.62	16.84	3.91	17.55	4.12	17.34	4.42	19.99	5.67	21.28
d-CoQA	ALL	12.50	37.80	13.92	40.18	17.01	43.27	14.82	39.92	20.85	49.09	20.97	49.35
	5000	-	-	12.86	38.27	14.09	39.52	13.21	38.88	18.30	45.49	17.63	44.94
	1000	-	-	11.45	35.43	12.63	38.92	12.39	38.11	15.43	42.75	16.08	42.27

Table 1: Results on d-QuAC and d-CoQA. Bold values indicate the best performance.

Model	EM	F1
BERT-S (Devlin et al. 2018)	1.81	32.62
UQACT (Lewis, Denoyer, and Riedel 2019)	5.32	42.23
CAsE (Cao et al. 2020)	5.75	44.28
AdaMRC (Wang et al. 2019)	7.20	43.11
RMRC-fix	9.25	52.08
RMRC	9.28	55.67

Table 2: Results on Alipay dataset.

Model	EM	F1
BERT-S	1.81	32.62
RMRC w/o r_b	8.44	54.13
RMRC w/o Answer Filtering	2.10	35.05
RMRC w/o Question Fusing	8.03	53.71
RMRC w/o QS Training	8.52	51.72
RMRC w/ Confidence-based Selector	8.76	52.34
RMRC w/ CE Reward	7.71	55.52
RMRC	9.28	55.67

Table 3: Performance of different variants of RMRC.

including two public datasets (QuAC (Choi et al. 2018) and CoQA (Reddy, Chen, and Manning 2019)) and one real-world dataset from Alipay. Note that, as the real-world Alipay data is in Chinese, we use another Chinese corpus instead of SQuAD for pre-training MRC model, i.e., a collection of CMRC (Cui et al. 2018), DRCD (Shao et al. 2018) and DUREADER (He et al. 2017).

QuAC and CoQA with Synthetic Noises: QuAC consists of 12,567 documents, including 11,567 training documents and 1,000 testing documents. Each document is affiliated with a related dialogue. In each round of chat in the dialogue, one user asks a question about the document, and the other user answers it. In total, there are 69,109 QA pairs for training and 5,868 for testing. Similar to QuAC, CoQA is composed of a training set of 7,199 documents along with 107,285 QA pairs, and a testing set of 500 documents along with 7,918 QA pairs. As the QA pairs of QuAC and CoQA are well-matched in the dialogue, we simulate the real conversation with noisy correspondence by randomly shuffling the questions in each dialogue. Specifically, for each dialogue, we randomly move each question ahead of the corresponding answer up to 1 ~ 5 rounds. We denote the shuffled datasets of QuAC and CoQA as d-QuAC and d-CoQA, re-

spectively in the following.

Alipay Dataset with Real Noises:¹ Alipay Dataset is collected in real-world scenarios, which contains the conversations about the marketing activities from the customer and customer service. In total, the dataset consists of 1,526 dialogues and 3,813 human-annotated QA pairs for testing.

Implementation

In our experiments, we take the widely-used BERT as the base encoder for the QS and the MRC encoder. The BERT network contains 12 hidden layers, each of which consists of 12 attention heads. The maximal input length and the hidden size are fixed to 512 and 768, respectively. The structure difference between QS and MRC is that the latter consists of two single-layered feed-forward networks for answer prediction. In the experiments, our MRC model is pre-trained on the SQuAD dataset (Rajpurkar et al. 2016) with the standard MRC training mechanism and then fine-tuned on the target domains (i.e., QuAC, CoQA). Similarly, the QS is also pre-trained on the SQuAD dataset by using the labeled QA pairs as positive samples and the randomly selected pairs as negative samples. Note that, as the Alipay data is in Chinese, we use another Chinese corpus instead of SQuAD for pre-training, i.e., a collection of CMRC (Cui et al. 2018), DRCD (Shao et al. 2018) and DUREADER (He et al. 2017).

For all experiments, we generate n-grams for each document by setting $K = 7$ for Eq. 1 and set the threshold for answer filtering γ to 0.7 and select the question by fixing κ in Eq. 6 to 5. The optimal parameters are determined by the grid search in the Alipay dataset and used for all experiments. We set the baseline score of the reward r_b to 0.7 for Eq. 14 in all experiments. For network training, we use the Adam optimizer whose learning rate is set to $2e - 5$ and $1e - 5$ for pre-training and fine-tuning, respectively. For evaluations, we take the *Exact Match* (EM) and *F1-score* (F1) as the performance measurements. Both the metrics are the higher the better.

Comparison Experiments

To show the effectiveness of our method, we compare it with the baselines including a vanilla BERT model trained in the source domain (denoted by BERT-S) and

¹The dataset is desensitized without any personal identifiable information.




Figure 3: (a-c) Influence of answer filtering threshold (γ), the selected question number (κ), and the baseline reward (r_b).

<p>Document: Tillis is known for his conservative stances on issues like same-sex marriage and women's abortion rights. He worked to nationalize the race by tying Hagan to President Obama, who suffers from high disapproval ratings in the Tar Heel state even though he -- along with Hagan -- were victorious there in 2008.</p>			
Extracted Answer	Corresponding Question	Previous Method	Our Method
same-sex marriage	What issues is he known for right-wing views on?	Where is Tillis for his conservative stances on issues like women's abortion rights ?	Epoch 1: And his rival? (-0.20) Epoch 2: Anything else? (-0.12) Epoch 3: What issues is he known for right-wing views on? (0.13) Epoch 4: What issues is he known for right-wing views on? (0.37)

<p>Document: Tuvalu has a population of 10,640 (2012 census). The total land area of the islands of Tuvalu is 26 square kilometres (10 sq mi). In 1568, Spanish navigator Álvaro de Mendaña was the first European to sail through the archipelago, sighting the island of Nui during his expedition in search of Terra Australis. In 1819 the island of Funafuti was named Ellice's Island; the name Ellice was applied to all nine islands after the work of English hydrographer Alexander George Findlay</p>			
Extracted Answer	Corresponding Question	Previous Method	Our Method
Funafuti	What island was named Ellice's island?	How much does the 1819 island of Funafuti cost Ellice's Island if any ?	Epoch 1: what are the Ellice islands now called? (-0.20) Epoch 2: Where is Tuvalu located? (-0.20) Epoch 3: Where is Tuvalu located? (-0.20) Epoch 4: what island was named Ellice's island? (0.80)

Figure 4: Comparisons between the questions constructed by our method and existing work (Lewis, Denoyer, and Riedel 2019). The red number and green number denote the negative and positive reward from MRC model, respectively.

three QA generation based MRC domain adaptation methods, i.e., UQACT (Lewis, Denoyer, and Riedel 2019), AdaMRC (Wang et al. 2019) and CAsE (Cao et al. 2020). Note that CAsE (Cao et al. 2020) uses the annotated questions in the target domain for fine-tuning, while the questions are unavailable in our experiment settings. As a remedy, we use a pre-trained question generator to generate questions in the target domain for CAsE. All the baselines and our method follow the same training pipeline, i.e., first pretraining the MRC model on the SQUAD data, then fine-tuning and evaluating it on the target datasets (d-QuAC, d-CoQA and Alipay). For all baselines, we conduct the experiments with the recommended parameters and report the best one. For our method, we repeat the experiments three times with different seeds and report the average. Moreover, we additionally use a variant of RMRC with a fixed QS (denoted as RMRC-fix) that keep the weights of QS fixed instead of optimizing with proposed reinforced self-training optimizer. Such a baseline could help us to understand the influence of

the proposed reinforced self-training strategy.

Results on d-QuAC and d-CoQA: Table 1 shows the quantitative results on d-QuAC and d-CoQA. From the results, one could observe that all domain adaptation methods outperform BERT-S. Moreover, the proposed methods (i.e., RMRC and RMRC-fix) significantly outperform the other three domain adaptation methods. Especially, RMRC achieves performance gains of **1.48/5.13** and **3.96/6.08** in d-QuAC-ALL and d-CoQA-ALL compared with the best baseline, respectively. From the comparisons between RMRC and RMRC-fix, one could easily find the effectiveness of our QS training strategy. Furthermore, to investigate the performance of the proposed method with small training datasets, we conduct experiments on QuAC and CoQA with 5,000 and 1,000 training samples. In other words, only 1.45% of original training QuAC samples are used in the evaluation with d-QuAC-1000. As shown, with the data size decreasing, the performance of all methods decreases while RMRC still outperforms all baselines.

Results on Alipay Dataset: Table 2 shows the quantitative results on the Alipay dataset. From the results, one could find that RMRC outperforms all baselines by a considerable performance margin, demonstrating its effectiveness in real-world dataset. Specifically, RMRC achieves performance gains of **2.08** and **12.56** over the best baseline (CAsE) in terms of EM and F1, respectively.

Ablation Study

In this section, we carry out ablation study on the Alipay dataset. In the experiments, we report the performance by removing or replacing some modules, including removing answer filtering (RMRC w/o Answer Filtering), question fusing (RMRC w/o Question Fusing), QS training (RMRC w/o QS Training); replacing f1-score reward by cross-entropy reward (RMRC w/ CE Reward), and reinforced self-training mechanism by confidence-based self-training mechanism (RMRC w/ Confidence-based Selector). As shown in Table 3, all the proposed modules are crucial to achieving encouraging results and the following conclusions could be obtained: i) All the RMRC variants outperform BERT-S, demonstrating the value of the dialogue-based QA pair construction; ii) RMRC and RMRC w/ CE Reward significantly outperform both RMRC w/o QS Training and RMRC w/ Confidence-based Selector, which shows the effectiveness of the proposed reinforced self-training algorithm for QS training; iii) RMRC that using F1 as the reward achieves better performance CE as F1 score directly measures the precision of the predicted answer.

Moreover, to further investigate the effectiveness of the proposed method without dialogues, we conduct experiments by only using the documents in QuAC and CoQA datasets. Specifically, we leverage the generators of UQACT (Lewis, Denoyer, and Riedel 2019), CAsE (Cao et al. 2020) and AdaMRC (Wang et al. 2019) to generate candidate questions from the given documents. Each generator generates one question. With the generated questions, we have two strategies: 1) keep all questions; 2) only select two candidate questions in terms of the relevance to the answers by the QS which will be optimized by our reinforced self-learning method. Then we construct the QA pairs similar to the experiments with dialogues and fine-tune the MRC model. The results are shown in Table 4. RMRC-all denotes using all generated questions, and RMRC-select denotes selecting the two most related questions by our QS model. As one could observe, with the proposed reinforced self-learning method, QS could select better questions for MRC fine-tuning thus achieving better performance.

Parameter Sensitivity Analysis

In this section, we carry out experiments on the Alipay dataset to investigate the influence of different parameters in RMRC including the answer filtering threshold γ in Eq. 3, the number of nearest questions κ in Eq. 6, and the baseline reward r_b in Eq. 14. As shown in Fig. 3(a), the performance keeps continuously increasing with increasing γ and gets the best result when $\gamma = 0.7$. As shown in Fig. 3(b), the performance of the MRC model keeps continuously increasing

Data	Size	CAsE		RMRC-all		RMRC-select	
		EM	F1	EM	F1	EM	F1
QuAC	ALL	6.79	19.33	6.45	20.38	6.80	21.54
	5000	4.85	17.92	4.83	18.82	5.77	20.92
	1000	3.91	17.55	4.22	18.20	4.57	19.89
CoQA	ALL	17.01	43.27	16.44	42.38	18.38	44.86
	5000	14.09	39.52	15.88	40.96	17.85	44.29
	1000	12.63	38.92	13.85	39.89	15.33	42.02

Table 4: Results on d-QuAC and d-CoQA without using dialogues. Bold values indicate the best performance.

with increasing κ from 1 to 5, showing that the correct question was more likely selected with more selected questions. As κ increases from 5 to 7, the performance significantly decreases. The reason is that too many selected questions inevitably introduce more noisy pairs. As shown in Fig. 3(c), the performance of the MRC model keeps improving when r_b increases from 0.1 to 0.7, and then decreases when r_b is larger than 0.7. The reason is that an over-high r_b will discourage almost all the predictions for QS training, thus leading to incorrect feedback for QS.

Case Study

To investigate the effectiveness of the QS and reinforced self-training optimizer, we conduct the case study in Fig. 4. As shown, the previous domain adaptation works will suffer from the NC problem, thus yielding irrelevant questions to the answer. In contrast, RMRC explicitly solves the NC problem by training QS with the MRC model feedback. Specifically, in the first epoch, QS would select an irrelevant question with a negative reward from the MRC model. After training with several epochs, the optimized QS could find the desirable question to the answer. Such an example intuitively demonstrates the effectiveness of our QS and the reinforced self-training algorithm.

Conclusion

This paper could be the first successful attempt to solve the noisy correspondence problem in the domain adaptation methods for MRC. Different from the well-studied noisy labels, the noisy correspondence refers to the errors in alignment instead of category-level annotation. To overcome this challenge, we propose a robust domain adaptation method for MRC with a novel reinforced self-training optimizer. Extensive experiments verify the effectiveness of the proposed method in leveraging synthesized and real-world dialogue data for MRC.

Acknowledgments

The authors would thank to the anonymous reviewers whose valuable suggestions and constructive comments remarkably improve this work. This work was supported in part by the National Key R&D Program of China under Grant 2020YFB1406702; in part by NSFC under Grant U21B2040

and 62176171; in part by Open Research Projects of Zhejiang Lab under Grant 2021KH0AB02.

References

- Arazo, E.; Ortego, D.; Albert, P.; O'Connor, N.; and McGuinness, K. 2019. Unsupervised label noise modeling and loss correction. In *International Conference on Machine Learning*, 312–321. PMLR.
- Bai, Y.; Yang, E.; Han, B.; Yang, Y.; Li, J.; Mao, Y.; Niu, G.; and Liu, T. 2021. Understanding and Improving Early Stopping for Learning with Noisy Labels. *Advances in Neural Information Processing Systems*, 34.
- Baradaran, R.; and Amirkhani, H. 2021. Ensemble learning-based approach for improving generalization capability of machine reading comprehension systems. *Neurocomputing*, 466: 229–242.
- Cao, Y.; Fang, M.; Yu, B.; and Zhou, J. T. 2020. Unsupervised domain adaptation on reading comprehension. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7480–7487.
- Choi, E.; He, H.; Iyyer, M.; Yatskar, M.; Yih, W.-t.; Choi, Y.; Liang, P.; and Zettlemoyer, L. 2018. Quac: Question answering in context. *arXiv preprint arXiv:1808.07036*.
- Cui, Y.; Liu, T.; Che, W.; Xiao, L.; Chen, Z.; Ma, W.; Wang, S.; and Hu, G. 2018. A span-extraction dataset for chinese machine reading comprehension. *arXiv preprint arXiv:1810.07366*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ding, Z.; Li, S.; Shao, M.; and Fu, Y. 2018. Graph adaptive knowledge transfer for unsupervised domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 37–52.
- Du, X.; Shao, J.; and Cardie, C. 2017. Learning to ask: Neural question generation for reading comprehension. *arXiv preprint arXiv:1705.00106*.
- Gao, S.; Ren, Z.; Zhao, Y.; Zhao, D.; Yin, D.; and Yan, R. 2019. Product-aware answer generation in e-commerce question-answering. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 429–437.
- Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I.; and Sugiyama, M. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *arXiv preprint arXiv:1804.06872*.
- He, W.; Liu, K.; Liu, J.; Lyu, Y.; Zhao, S.; Xiao, X.; Liu, Y.; Wang, Y.; Wu, H.; She, Q.; et al. 2017. Dureader: a chinese machine reading comprehension dataset from real-world applications. *arXiv preprint arXiv:1711.05073*.
- Hu, L.; Kan, M.; Shan, S.; and Chen, X. 2018. Duplex generative adversarial network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1498–1507.
- Huang, Z.; Niu, G.; Liu, X.; Ding, W.; Xiao, X.; Peng, X.; et al. 2021. Learning with Noisy Correspondence for Cross-modal Matching. In *Advances in Neural Information Processing Systems*.
- Jiang, Y.; and Zhao, Z. 2018. StackReader: An RNN-Free Reading Comprehension Model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Kan, M.; Shan, S.; and Chen, X. 2015. Bi-shifting auto-encoder for unsupervised domain adaptation. In *Proceedings of the IEEE international conference on computer vision*, 3846–3854.
- Lewis, P.; Denoyer, L.; and Riedel, S. 2019. Unsupervised question answering by cloze translation. *arXiv preprint arXiv:1906.04980*.
- Li, H.; Zhang, X.; Liu, Y.; Zhang, Y.; Wang, Q.; Zhou, X.; Liu, J.; Wu, H.; and Wang, H. 2019. D-NET: A simple framework for improving the generalization of machine reading comprehension. In *EMNLP 2019 MRQA Workshop*, 212.
- Li, W.; Duan, L.; Xu, D.; and Tsang, I. W. 2013. Learning with augmented features for supervised and semi-supervised heterogeneous domain adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 36(6): 1134–1148.
- Liu, T.; and Tao, D. 2015. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*, 38(3): 447–461.
- Luo, Y.; Han, B.; and Gong, C. 2021. A bi-level formulation for label noise learning with spectral cluster discovery. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, 2605–2611.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Reddy, S.; Chen, D.; and Manning, C. D. 2019. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7: 249–266.
- Seo, M.; Kembhavi, A.; Farhadi, A.; and Hajishirzi, H. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- Shao, C. C.; Liu, T.; Lai, Y.; Tseng, Y.; and Tsai, S. 2018. Dred: a chinese machine reading comprehension dataset. *arXiv preprint arXiv:1806.00920*.
- Song, H.; Kim, M.; Park, D.; Shin, Y.; and Lee, J.-G. 2020. Learning from noisy labels with deep neural networks: A survey. *arXiv preprint arXiv:2007.08199*.
- Su, D.; Xu, Y.; Winata, G. I.; Xu, P.; Kim, H.; Liu, Z.; and Fung, P. 2019. Generalizing question answering system with pre-trained language model fine-tuning. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, 203–211.
- Wang, H.; Gan, Z.; Liu, X.; Liu, J.; Gao, J.; and Wang, H. 2019. Adversarial domain adaptation for machine reading comprehension. *arXiv preprint arXiv:1908.09209*.

- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3): 229–256.
- Xia, X.; Liu, T.; Han, B.; Gong, C.; Wang, N.; Ge, Z.; and Chang, Y. 2020. Robust early-learning: Hindering the memorization of noisy labels. In *International Conference on Learning Representations*.
- Yang, M.; Huang, Z.; Peng, H.; Li, T.; Lv, J. C.; and Peng, X. 2022. Learning with Twin Noisy Labels for Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yang, M.; Li, Y.; Huang, Z.; Liu, Z.; Hu, P.; and Peng, X. 2021. Partially view-aligned representation learning with noise-robust contrastive loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1134–1143.
- Yu, A. W.; Dohan, D.; Luong, M.-T.; Zhao, R.; Chen, K.; Norouzi, M.; and Le, Q. V. 2018. Qanet: Combining local convolution with global self-attention for reading comprehension. *arXiv preprint arXiv:1804.09541*.
- Yu, X.; Han, B.; Yao, J.; Niu, G.; Tsang, I.; and Sugiyama, M. 2019. How does disagreement help generalization against label corruption? In *International Conference on Machine Learning*, 7164–7173. PMLR.
- Zhang, C.; Ding, H.; Lin, G.; Li, R.; Wang, C.; and Shen, C. 2021. Meta navigator: Search for a good adaptation policy for few-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9435–9444.
- Zhang, W.; Ouyang, W.; Li, W.; and Xu, D. 2018. Collaborative and adversarial network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3801–3809.
- Zhang, Z.; Yang, J.; and Zhao, H. 2020. Retrospective reader for machine reading comprehension. *arXiv preprint arXiv:2001.09694*.