

Multi-source Machine Reading Comprehension with Meta-Learning and Adaptive Adversarial Training

Zhe Wen^{*†}, Yatao Qi^{*†}, Huan Liu^{*✉}, Zheng Lin^{*†}, Weiping Wang^{*}

^{*}Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China

[†]School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China

{wenzhe, qiyatao, liuhuan, linzheng, wangweiping}@iie.ac.cn

Abstract—Machine reading comprehension (MRC) methods have shown great success in many datasets, but existing methods fail to achieve satisfactory results in low-resource scenarios. In addition, existing MRC models suffer from a notable decrease in performance when confronted with scenes different from the training data. Thus, it is hard to transfer knowledge between domains. In this paper, we propose an adaptive meta-learning framework to learn and transfer the shared knowledge. The framework is based on model-agnostic meta-learning algorithm, aiming to aggregate meta-knowledge among multi-source datasets. Furthermore, for better adaptation to different target domain, we investigate an adaptive adversarial training strategy to obtain domain-specific meta-knowledge. We empirically adopt three large-size datasets as source domains and five small-size datasets as target domains, and extensive experiments show the effectiveness of our framework.

Index Terms—Natural Language Understanding, Machine Reading Comprehension, Meta Learning, Adversarial Training.

I. INTRODUCTION

Machine reading comprehension (MRC) attempts to answer many kinds of questions, including simple questions and multi-hop reasoning questions, after reading a document or multiple documents. Recently, MRC models based on the manner of fine-tuning the pre-trained language model for downstream tasks, have achieved acceptable performance [1]. However, these models still rely on an adequate scale of the training data in the target domain. When confronted with the low-resource MRC task, even the performance of models based on pre-training deteriorates significantly.

As shown in Figure 1, as the scale of the labeled data reduced, the performance of BERT-based MRC model decreases by a large margin. Specifically, when approximately 40% data is used for training, the performance curve falls sharply. In fact, for some domains, the lack of adequate labeled data makes it difficult to train a domain-specific MRC model with satisfactory performance. Therefore, it is of great importance and necessity to solve the low-resource problem in MRC.

An intuitive solution is to generalize or transfer knowledge learned in resource-rich domains to resource-poor domains. [2] trained models on one or more source MRC datasets, and then evaluated their performance on a target domain, either without any additional target training examples (generalization) or with additional target examples (transfer). They found that current

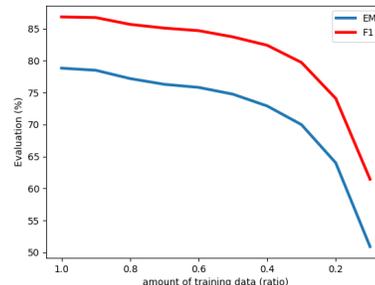


Fig. 1: Exact match and F1 curves of the BERT-based MRC model on SQuAD, with different data volume.

models over-fit to the particular training set and generalize poorly even to similar datasets. Besides, they also investigated the effect of knowledge transfer based on domain-level similarities and fine-tuning. Despite some improvements, the shared knowledge among multiple datasets of source domains has not been fully exploited. Although each MRC dataset involves the use of different expertise, there is still substantial overlap in the abilities required to answer questions across these datasets. Therefore, we believe it is meaningful to effectively exploit the shared knowledge among multiple datasets, and use them to promote the language understanding of the low-resource setup.

For a target domain, the learning objective is to find an optimal model that is well adapted to the current data. However, due to the lack of sufficient training samples, the optimal parameters oscillate with time during training. Therefore, it is difficult to learn an acceptable model with insufficient data. To address this issue, we propose an adaptive meta-learning framework for multi-source domain adaptation of MRC. The intuition behind the framework is that the past experience over multiple source domains can be used to facilitate the adaptation to a new environment. Specifically, we employ a model-agnostic meta-learning (MAML) algorithm [3] to aggregate the meta-knowledge from multiple source domains. The main focus of this paper is not to propose a new meta-learning algorithm, but to explore whether the representative meta-learning algorithm MAML is suitable for many-to-one machine reading comprehension scenario, and how to improve the domain transfer ability under the MAML framework.

Since meta-knowledge is transferred, there is still a gap between source domains and target domains. Especially when

✉ Corresponding author.

DOI reference number: 10.18293/SEKE2023-179.

the difference of data distribution is large, the performance of domain adaptation will be highly affected. To fill the domain gap, one possible way is to generate adversarial examples through adding imperceptible perturbations on inputs, which can be used to improve the model robustness [4]. However, the traditional adversarial training adds noises with the gradient of the source domain, which is futile to cross the domain gap. It may even broaden the gap. To solve this issue, we design an adaptive meta-learning framework based on adaptive adversarial training strategy, and employ gradient of the target domain to guide adversarial examples generation in the source domain. Therefore, we can obtain domain-specific meta-knowledge, rather than the fixed one.

The main contributions can be summarized as follows:

- To overcome the challenge of low-resource MRC, we propose an adaptive meta-learning framework. With this framework, the meta knowledge across multiple source domains can be learned, which is further used to enhance the target-domain MRC ability. Besides, it can be easily combined with arbitrary MRC models like BERT.
- To bridge the gap between domains in meta-learning, we devise an adaptive adversarial training strategy. It helps obtain the domain-specific meta knowledge. For both the generality and the individuality are taken into consideration, it is more suitable for each target domain.
- We evaluate our model on three resource-rich source datasets and five resource-poor target datasets, and the results demonstrate that the proposed model consistently outperforms state-of-the-art baselines considerably.

II. METHODOLOGY

A. Task Definition

In general, MRC tasks can be categorized into the extractive setting and the multiple-choice setting. To facilitate knowledge transfer, we formulate all the MRC datasets to the extractive type. Therefore, the example in all datasets contains a question, a paragraph, and an answer whose content is a local context extracted from the original paragraph.

Formally, given the paragraph P and a question Q , the goal of MRC is to extract an answer span $[s, e]$ from P , where s denotes the start position and e denotes the end position in P .

B. Framework of Adaptive Meta-Learning

Our adaptive meta-learning framework consists of two parts: Multi-source Meta-Learning (MML) and Adaptive Adversarial Training (AAT). In the process of MML, the meta-knowledge is learned on the merged data of multiple source domains $D^S = \{D_1^S, D_2^S, \dots, D_n^S\}$ through MAML [3].

In the process of AAT, the adaptive adversarial examples are generated, which can be used in MML to encourage the meta-model to learn toward the target domain. Specifically, the generation of adaptive adversarial samples $\tilde{D}_{T_i}^S$ is guided by the gradient of the target domain $\{D_1^T, D_2^T, \dots, D_m^T\}$, respectively. Then, $\tilde{D}_{T_i}^S$ participates in the meta-learning process together with D^S in a data-augmentation way. The details of adaptive meta-learning are summarized in Algorithm 1.

Algorithm 1 The procedure of Adaptive Meta-Learning

Input: Multi-source data D^S , target data D_i^T , three sets of parameters θ^* , $\hat{\theta}$, θ_t , and their learning rate α , β , γ

Output: Optimized parameters θ^*

- 1: Initialize θ^* , $\hat{\theta}$, θ_t ;
 - 2: Fine-tune the target model parameters θ_t on data D_i^T :
 $\theta_t' = \theta_t - \gamma \nabla_{\theta_t} L(f(\theta_t, D_i^T))$;
 - 3: **for** all X^S from D^S **do**
 - 4: Compute gradients on θ_t with data X^S , and generate adversarial samples:
 $\tilde{X}^S = X^S + \lambda * \nabla_{\theta_t}$
 - 5: **end for**
 - 6: Merge \tilde{D}^S and D^S , get new source data $D^{S'}$;
 - 7: **while** not convergent **do**
 - 8: Sample data $\{D^{S_1}, \dots, D^{S_K}\}$ from $D^{S'}$;
 - 9: **for** $i = 1$ to K **do**
 - 10: Sample support set $D_{Sup}^{S_i}$ from D^{S_i} and train the model $\hat{\theta}$:
 $\hat{\theta}' = \hat{\theta} - \alpha \nabla_{\hat{\theta}} L(f(\hat{\theta}, D_{Sup}^{S_i}))$
 - 11: Sample query set $D_{Que}^{S_i}$ from D^{S_i} and compute gradients for accumulation:
 $\theta_{acc+} = \hat{\theta}$, $\nabla_{\theta_{acc+}} = \nabla_{\hat{\theta}} L(f(\hat{\theta}, D_{Que}^{S_i}))$
 - 12: **end for**
 - 13: Meta model update:
 $\theta^* = \theta_{acc}/K$, $\theta^{*'} = \theta^* - \beta \nabla_{\theta_{acc}}/K$;
 - 14: **end while**
 - 15: Transfer θ^* as the initial parameters to the downstream task D_i^T and fine-tune.
-

C. Multi-source Meta-Learning

MML is used to aggregate the meta-knowledge from multiple resources. In the meta-learning problem setting, the goal is to learn models that can learn new tasks from small amounts of data. To achieve this, meta-learning algorithms require a set of meta-training and meta-testing tasks drawn from a certain distribution. The key assumption of learning-to-learn is that the tasks in this distribution share common structure that can be exploited for faster learning of new tasks.

There are two sets of model parameters required in meta-learning. One set of parameters is used for the meta-training task denoted as $\hat{\theta}$, another set is used for the meta-testing task denoted as θ^* . For a downstream task, θ^* is employed for model initialization and to be fine tuned afterward. The process of MML is described as follows:

Meta-learning involves multiple sampling data subsets $D^{Sample} = \{D^{S_1}, D^{S_2}, \dots, D^{S_K}\}$. Each sampling subset D^{S_i} includes a support set and a query set, denoted as $D^{S_i} = \{D_{Sup}^{S_i}, D_{Que}^{S_i}\}$, and each pair of the support set and the query set is sampled from the same distribution.

Fast Optimization In a sampling subset, we utilize the support set $D_{Sup}^{S_i}$ to optimize $\hat{\theta}$. This process can be regarded as applying $D_{Sup}^{S_i}$ as the training data to train a model, and the model parameters are optimized by stochastic gradient descent.

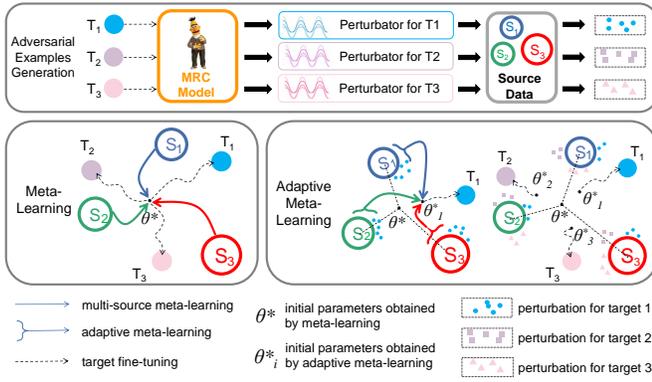


Fig. 2: Overview of the adaptive meta-learning framework. The larger circles with solid line represent source domains, and the smaller brimless circles represent target domains. Adaptive meta-learning is conducted by adding adaptive adversarial examples into multi-source meta-learning.

$$\hat{\theta}' = \hat{\theta} - \alpha \nabla_{\hat{\theta}} L_{D_{Sup}^S} (f(\hat{\theta})), \quad (1)$$

where L denotes the cross-entropy loss function, and α denotes the learning rate.

Gradients Calculation After the previous step, we can obtain a set of optimized parameters. Then, we utilize these parameters to calculate the gradients on the respective query set. This process can be regarded as applying D_{Que}^S as the development data to evaluate the sampled task.

Gradients Accumulation Next, the gradients of multiple sampled data points are accumulated. Since the accumulation operation considers the gradient direction generated from various subsets, it can be regarded as the shared knowledge across multi-domain being learned.

Parameters Update The cumulative gradient is used to update the parameters of the meta-model. Each update of θ^* is seen as one step of training for meta-model. The formula used to update the parameters is as follows:

$$\theta^{*'} = \theta^* - \beta \nabla_{\hat{\theta}_{acc}} / K, \quad (2)$$

where $\hat{\theta}_{acc}$ denotes the accumulative gradients on the parameters $\hat{\theta}$, and β is the learning rate, and K is the number of sampling tasks.

Target Fine-tuning Through multiple iterations of the above steps, a set of meta-model parameters θ^* is learned. Then, we fine-tune the target model with the initial parameters θ^* and the target dataset D_i^T , to enable the model to be close to the target domain.

D. Adaptive Adversarial Training

Under the framework of meta-learning, we further propose Adaptive Adversarial Training (AAT) aiming to bridge the gap between domains. It is based on the idea of improving model generalization capability by adversarial training. Adversarial training uses a mixture of adversarial and original examples

to train a model to make it more robust. Here, the adversarial examples are in the form of continuous vectors instead of actual texts in natural languages. Specifically, adversarial examples are slightly different from original examples, and they are generated by adding imperceptible perturbation to the embeddings of the original examples. We first feed the source data D^S into a frozen well-trained target domain network, and then generate adversarial samples $\tilde{D}_{T_i}^S$ in the direction of the target domain. The process of AAT is described as follows:

Perturbator Training To generate adversarial samples more adapted to the target domain, we take the target model as a perturbator. After fine-tuning on the target data D_i^T , we acquire a target model θ_t , which is used to generate directional perturbations. The perturbator is optimized as:

$$\theta_t' = \theta_t - \gamma \nabla_{\theta_t} L(f(\theta_t, D_i^T)), \quad (3)$$

where γ is the learning rate and L denotes the cross-entropy loss function in MRC

Adversarial Examples Generation Next, the perturbator θ_t and the multi-source data D^S are used to generate the adaptive adversarial examples. Specifically, the objective of this step is to increase the loss of the target model, given X_s sampled from D^S and its added perturbation generated from D_i^T . Through multiple iterations of the above step, we can obtain the set of adaptive adversarial examples $\tilde{D}_{T_i}^S$ corresponding to D^S . Then, $\tilde{D}_{T_i}^S$ is used to train the meta-model in MML, as well as the previous source-domain data.

III. EXPERIMENTS

A. Datasets

In the experiments, we choose eight reading comprehension datasets with different sizes. According to the size of each dataset, we briefly divide the datasets into source domain and target domain. Specifically, three larger ones SQuAD [5], NewsQA [6] and HotpotQA [7] are merged into multi-source data, and five smaller ones DROP [8], WikiHop [7], ComQA [9], ComplexQuestion [10] and BoolQ [11] are regarded as target dataset respectively.

Our adaptive meta-learning method transfers the meta-knowledge aggregated from multi-source domain datasets to each target dataset. Due to the insufficient computation resources, we randomly sample 100K examples from each source dataset and combine them as the whole source-domain dataset.

B. Baselines

We mainly compare our method with several popular pre-trained models, which present state-of-the-art performance in recent MRC tasks. These pre-trained models are directly fine-tuned on target dataset. Besides, we also compare our method with the single-source transfer method and the multi-source transfer method. For all methods, we adopt the Exact Match (EM) and the F1 score as our evaluation metrics.

BERT, XLM, XLNet, RoBERTa: They are Transformer-based pre-trained language models.

SynNet: SynNet [12] is a data augmentation method using a two-stage synthesis network for MRC.

IMM: IMM [13] is an iterative multi-source mutual knowledge transfer framework for MRC.

MADE: MADE [14] is to model multi-dataset question answering with an ensemble of single-dataset experts, by training a collection of dataset-specific adapter modules that share an underlying Transformer model.

SingleSource-T: We use BERT to train a MRC model on single-source dataset, and then transfer the trained model to target domain and fine-tune it with target dataset.

MultiQA: MultiQA [2] is a BERT based MRC model, which is trained on merged multi-source dataset, and then transfer the trained model to the target domain.

SingleSource-M: We use MAML to learn meta knowledge on single-source dataset, and then transfer the meta knowledge to target domain and retrain it with target dataset.

MultiQA-M: We use MAML to learn meta knowledge on merged multi-source dataset, and then transfer the meta knowledge to target domain and retrain it with target dataset.

C. Experimental Setup

In the experiments, for each target domain and each baseline model, only 300 labeled samples are used for domain adaptation. For the MRC task, it can be seen as a low-resource scenario, because a few hundreds of samples are insufficient to learn a robust model. In the MADE [14], all the 300 examples are used to fine-tune the adapters. In the meta-learning, the 300 samples are divided into support set and query set. The number of the support set and the query set are set to 280 and 20 respectively. We implement our model based on the *Transformers* released by Hugging Face, running on a TITAN RTX GPU. The reported experimental results of our model are average of 5 runs. Following the settings of MultiQA [2], we set the training batch size to 24 and three learning rate α, β, γ uniformly to $3e-5$. We set maximum length of input sequence to 384. The length of some documents exceeds the limitation, thus we tailor the original document into pieces with the sliding window size of 384 and sliding stride of 128. All models adopt Adam with $\beta_1 = 0.9$, $\beta_2 = 0.999$ for optimization and generally converge after 2 epochs.

D. Overall Performance

In order to evaluate the effectiveness of the proposed method, we test various baselines on five target datasets. Table I presents the main experimental results. From Table I, we can make three important observations.

First, the models whose knowledge transferred from multiple source datasets perform better than those transferred from the single dataset in most cases. For the BERT-based (non-meta) knowledge transfer methods, MultiQA performs consistently better than NewsQA-T, HotpotQA-T and SQuAD-T, on EM and F1 respectively. For the meta knowledge transfer methods, MultiQA-M achieves a notable performance gain against NewsQA-M, HotpotQA-M and SQuAD-M. This is in accordance with our intuition, which indicates that there are

TABLE I: Exact Match and F1 Score on Five Target Datasets with Different Methods. For Each Target Dataset and Each Evaluation Metric.

Model	Drop		WikiHop		ComQA		ComplexQ		BoolQ	
	EM	F1	EM	F1	EM	F1	EM	F1	EM	F1
BERT	46.51	54.04	19.29	22.27	12.25	14.37	23.31	27.41	64.83	65.75
XLM	48.16	57.94	16.20	21.22	12.96	15.76	28.22	37.46	66.64	67.47
XLnet	44.00	53.20	16.14	19.16	8.24	10.02	30.06	38.18	64.50	64.50
Roberta	48.05	56.51	17.04	20.11	10.01	11.95	15.34	20.74	62.17	62.17
NewsQA-T	46.92	55.06	20.86	24.72	21.44	27.30	39.26	44.54	67.13	67.80
HotpotQA-T	47.19	55.24	21.88	26.01	22.03	27.80	39.26	43.68	67.40	68.18
SQuAD-T	48.54	56.48	21.33	25.41	27.56	34.38	38.65	44.74	66.94	67.27
MultiQA	51.07	58.17	21.94	26.30	30.98	37.48	42.94	49.85	69.45	69.92
SynNet	37.03	45.82	19.98	21.71	29.75	31.23	40.38	46.83	64.08	64.14
IMM	40.29	48.73	20.00	23.82	23.20	29.56	34.97	43.32	64.34	64.51
MADE	48.65	56.69	31.06	39.12	43.62	52.13	46.20	55.38	66.36	66.69
NewsQA-M	47.91	55.93	21.49	25.91	21.08	26.07	34.97	41.51	64.56	66.21
HotpotQA-M	47.52	55.40	21.78	25.79	24.26	29.88	34.97	40.23	69.39	69.78
SQuAD-M	48.65	56.09	23.31	27.64	28.15	35.23	39.26	47.58	60.98	66.31
MultiQA-M	52.99	60.34	25.45	29.98	33.45	41.10	43.56	50.95	69.45	69.92
Ours	59.68	65.92	25.67	30.15	36.98	46.51	44.17	52.03	73.06	74.28

shared abilities required to answer questions across different datasets, and exploring the shared abilities from multiple source datasets instead of the single one is helpful to improve the performance of the target-domain MRC.

Second, MultiQA-M performs better than the fine-tuned methods, including BERT, XLM [15], XLNet [16] and Roberta [17], and also outperforms the traditional knowledge transfer methods, including SynNet, IMM and MultiQA. The reason for the unsatisfactory results of fine-tuned methods mainly lies in the lack of sufficient training data for the target domain, and high-capacity parameters of the pre-training models. In general, pre-trained representations improve generalization, but their effect is moderate when the training data is insufficient. Besides, due to the meta-knowledge over multiple source domains promoting the fast convergence and adaptation to the target domain, transferring meta-knowledge is superior to the traditional knowledge. It suggests that meta-learning over multiple source domains is good at capturing the shared knowledge regardless of the data imbalance and different focus.

Third, our model achieves the best performance on all target datasets, and outperforms the baselines by a large margin. It indicates that not only the meta-knowledge but also the domain-specific meta-knowledge contribute to the effectiveness of our model. From the experimental results, we can infer that, in view of the meta-knowledge, there is still a gap between source domains and target domains. And correspondingly, the proposed adaptive meta-learning strategy is effective to fill the gap. Because different meta-knowledge can be obtained in the direction of different target domain under our design, the domain-specific meta-knowledge is well adapted to each target domain. In general, our design takes both generality and individuality into account, which is more appropriate for multi-source domain adaptation in MRC.

As shown in Table I, our model and MADE model have different strengths on the target domain datasets. In fact, the two models follow two different technique lines, thus having

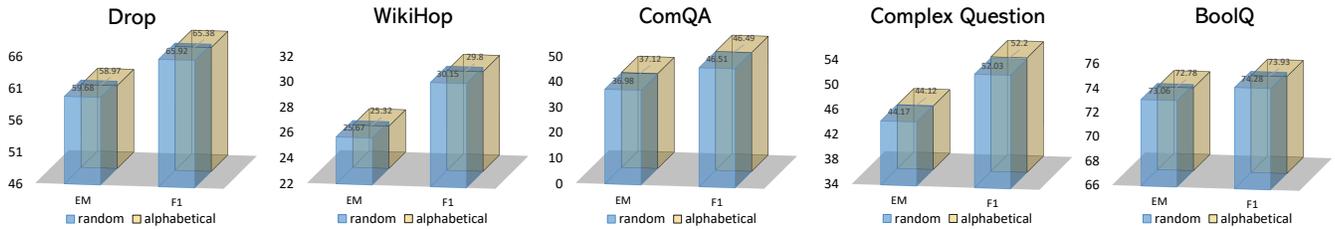


Fig. 3: Performance on randomly sampling 100K data and alphabetically selected 100K data, respectively.

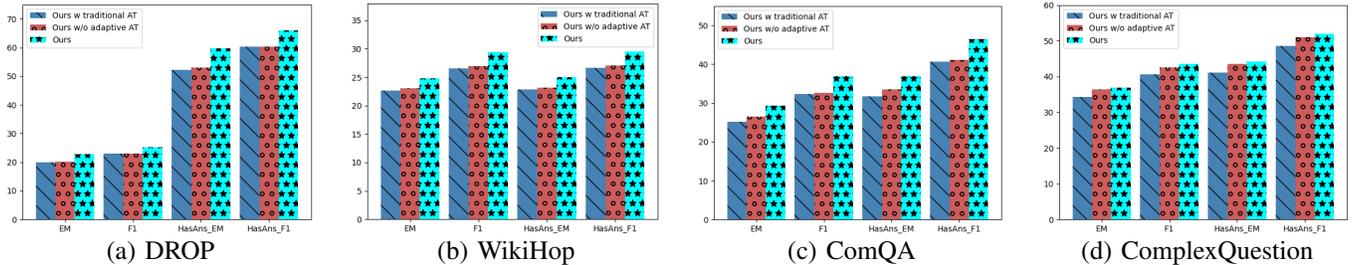


Fig. 4: Ablation study of our framework on various datasets. For each target dataset, two groups of EM and F1 are given. One group is evaluated on all examples, and the other group is evaluated on the examples containing the answers.

different expertise. Specifically, MADE is based on adapters, which provide a parameter-efficient alternative for the full fine-tuning in which we can only fine-tune lightweight neural network layers on top of pre-trained weights. However, our method is based on meta-learning, which helps optimize parameters to give better predictions in shorter time. Though the performance of two models is almost comparable, our method has two advantages. First, our method is model-agnostic and can be easily equipped with various machine learning models, including but not limited to pre-trained language models. Second, our method can train a more generalized model, which facilitates the learning of a novel task. For example, MADE performs better on WikiHop, ComQA and ComplexQuestion, but performs dramatically worse on Drop and BoolQ. Maybe it is because the former ones are all collected from Wikipedia, resulting in the similar question answering process, whereas Drop is crowd-sourced and adversarially created, and BoolQ provides quite different question type. By contrast, our method performs relatively stable among these datasets.

Besides, it is worth noting that the results in Table I are lower than those reported in the Leaderboards. The reasons lie in two aspects. First and foremost, the goal of our work is not to design a highly complicated model or pre-processing technique to significantly surpass the existing methods. Instead, we hope to explore a direction for MRC domain adaptation and other similar transfer learning problems. In fact, some of the top ranked models on the Leaderboards are based on super large-scale language models, e.g., T5 and GPT3. Due to the limitation of computational capacity, we finally adopt the widely used Bert-base as our backbone model, but the proposed cross-domain method can be easily combined with the existing models. Second, the datasets used for the manuscript and the Leaderboards are different. In our work,

for each target domain, only 300 labeled samples are used for domain adaptation, while for the Leaderboards all labeled samples could be used.

To better measure the impact of randomness on the dataset, we have additionally compared the randomly sampling 100K data with intentionally selected 100K data from the multi-source dataset. Specifically, the intentionally selected 100K samples are ranked alphabetically. Figure 3 presents the comparison results. As seen from Figure 3, our model performs basically equal on two kinds of datasets. It proves that our proposed method is less susceptible to the randomness on the quality of the combined multi-source dataset, to some extent.

E. Ablation Study

The main highlight of our work is the application of adaptive adversarial training in meta-learning for multi-source domain adaptation of MRC. To evaluate the effect of adaptive adversarial training, we compare our framework with its following variants: (1) Meta-learning without adversarial training. (2) Meta-learning with traditional adversarial training, i.e., the perturbations are added in the direction of the source domain. (3) Meta-learning with adaptive adversarial training, i.e., the perturbations are added in the direction of the target domain.

Owing to the limitation of space, Figure 4 only illustrates the comparison results on four target datasets. From Figure 4, we can observe that compared with general meta-learning, the introduction of traditional adversarial training does not improve performance. In fact, the effectiveness of adversarial training mainly owns to its ability of local smoothness around data points, and pushing decision boundary away from data points. However, the generated examples surround the original examples, which has only limited ability to fill the domain gap. In order to guide the direction of generating adversarial examples from source domains towards target domain,

we further propose adaptive adversarial training. Moreover, traditional adversarial training under transfer scenario may further widen the gap between the source domain and target domain, thus resulting in a performance drop. Overall, it is surprising to find that meta-learning with adaptive adversarial training significantly outperforms the meta-learning with or without traditional adversarial training. The reason is that the proposed adaptive adversarial training can guide the generation of adversarial examples to the direction of the target domain, then the model trained with these adversarial examples can boost the adaptation ability.

IV. RELATED WORK

a) Transfer learning for MRC: MRC is a challenging Natural Language Processing research field with wide applications [18]. To address the low-resource MRC problem, previous studies usually use pre-trained language models to enhance the performance of MRC models, such as BERT. Meanwhile, many efforts have been made to use transfer learning methods to improve MRC performance in target domain via datasets in source domain, among which fine-tuning is the most common technique [19].

Recently, [20] presented an alternate meta-learning framework for complex question answering over knowledge bases, which jointly and alternately optimized a retriever to select questions, and a programmer to produce an answer. However, this framework is designed for the scenario of question distributional bias within the same dataset instead of learning shared knowledge among different datasets. Differently, we aim to learn and transfer the shared knowledge among multiple MRC datasets through adaptive meta-learning. MADE [14] and our work share the similar motivation, but lie in different technique lines. MADE is based on adapter, which can be used as a parameter-efficient alternative to fine-tuning. We provide an investigation of meta learning combined with adaptive adversarial training, which is a missing part in the literature.

b) Meta-Learning: Meta-learning exploits inherent structures in data to learn more effective learning rules for fast domain adaptation [21]. To make the model sensitive to the new task, one popular direction of meta-learning is to train a meta-learner to learn how to update the parameters of the target model. [3] proposed Model Agnostic Meta-Learning (MAML) that constrains the learning rule for the model and uses stochastic gradient descent to quickly adapt networks to new tasks. This pattern has been widely used to adapt deep networks to resource-poor environments. S2A [22] first trained a retriever and a programmer separately, and then established a meta-learning task for questions and employed MAML to fine-tune the programmer, based on the similar samples found by the retriever. Unlike the above works, we apply MAML to the multi-domain MRC task with an adaptive adversarial training in its framework, and it encourages the learned knowledge to fit each target domain well.

V. CONCLUSION

In this work, we propose an adaptive meta-learning framework for multi-source domain adaptation of MRC. With this framework, we can learn the meta-knowledge, which implies the shared abilities of reading comprehension, across multiple source datasets. Furthermore, an adaptive adversarial training strategy is devised to obtain domain-specific meta-knowledge, which fills the gap between source and target domains. Extensive experiments demonstrate that our framework establishes a state-of-the-art baseline for low-resource MRC.

REFERENCES

- [1] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [2] A. Talmor and J. Berant, "Multiqa: An empirical investigation of generalization and transfer in reading comprehension," in *ACL*, 2019.
- [3] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *ICML*, 2017, pp. 1126–1135.
- [4] P. Jiang, A. Wu, Y. Han, Y. Shao, M. Qi, and B. Li, "Bidirectional adversarial training for semi-supervised domain adaptation," in *IJCAI*, 2020, pp. 934–940.
- [5] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang, "Squad: 100,000+ questions for machine comprehension of text," *arXiv preprint arXiv:1606.05250*, 2016.
- [6] A. Trischler, T. Wang, X. Yuan, J. Harris, A. Sordoni, P. Bachman, and K. Suleman, "Newsqa: A machine comprehension dataset," *arXiv preprint arXiv:1611.09830*, 2016.
- [7] J. Welbl, P. Stenetorp, and S. Riedel, "Constructing datasets for multi-hop reading comprehension across documents," *TACL*, 2018.
- [8] D. Dua, Y. Wang, P. Dasigi, G. Stanovsky, S. Singh, and M. Gardner, "Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs," *arXiv preprint arXiv:1903.00161*, 2019.
- [9] A. Abujabal, R. S. Roy, M. Yahya, and G. Weikum, "Comqa: A community-sourced dataset for complex factoid question answering with paraphrase clusters," *arXiv preprint arXiv:1809.09528*, 2018.
- [10] A. Talmor and J. Berant, "Partitioning of the complexwebquestions dataset," *arXiv preprint arXiv:1807.09623*, 2018.
- [11] C. Clark, K. Lee, M.-W. Chang, T. Kwiatkowski, M. Collins, and K. Toutanova, "Boolq: Exploring the surprising difficulty of natural yes/no questions," in *Proceedings of NAACL-HLT*, 2019, pp. 2924–2936.
- [12] D. Golub, P.-S. Huang, X. He, and L. Deng, "Two-stage synthesis networks for transfer learning in machine comprehension," *arXiv e-prints*, pp. arXiv-1706, 2017.
- [13] X. Liu, K. Liu, X. Li, J. Su, Y. Ge, B. Wang, and J. Luo, "An iterative multi-source mutual knowledge transfer framework for machine reading comprehension," in *IJCAI*, 2020, pp. 3794–3800.
- [14] D. Friedman, B. Dodge, and D. Chen, "Single-dataset experts for multi-dataset question answering," *arXiv preprint arXiv:2109.13880*, 2021.
- [15] A. Conneau and G. Lample, "Cross-lingual language model pretraining," *Advances in neural information processing systems*, vol. 32, 2019.
- [16] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, "Xlnet: Generalized autoregressive pretraining for language understanding," *NeurIPS*, vol. 32, 2019.
- [17] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [18] J. Liu, M. Yu, Y. Chen, and J. Xu, "Cross-domain slot filling as machine reading comprehension: A new perspective," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 673–685, 2022.
- [19] Y.-A. Chung, H.-Y. Lee, and J. Glass, "Supervised and unsupervised transfer learning for question answering," in *NAACL-HLT*, 2018.
- [20] Y. Hua, Y.-F. Li, G. Haffari, G. Qi, and W. Wu, "Retrieve, program, repeat: complex knowledge base question answering via alternate meta-learning," *arXiv preprint arXiv:2010.15875*, 2020.
- [21] S. Li, X. Wang, Y. Cao, F. Xue, Z. Yan, and H. Zha, "Self-supervised deep visual odometry with online adaptation," in *CVPR*, 2020.
- [22] D. Guo, D. Tang, N. Duan, M. Zhou, and J. Yin, "Coupling retrieval and meta-learning for context-dependent semantic parsing," in *ACL*, 2019.