

Unsupervised Domain Adaptation of Language Models for Reading Comprehension

Kosuke Nishida, Kyosuke Nishida, Itsumi Saito, Hisako Asano, Junji Tomita

NTT Media Intelligence Laboratories, NTT Corporation

1-1 Hikarinooka Yokosuka-Shi, Kanagawa, Japan

kosuke.nishida.ap@hco.ntt.co.jp

Abstract

This study tackles unsupervised domain adaptation of reading comprehension (UDARC). Reading comprehension (RC) is a task to learn the capability for question answering with textual sources. State-of-the-art models on RC still do not have general linguistic intelligence; i.e., their accuracy worsens for out-domain datasets that are not used in the training. We hypothesize that this discrepancy is caused by a lack of the language modeling (LM) capability for the out-domain. The UDARC task allows models to use supervised RC training data in the source domain and only unlabeled passages in the target domain. To solve the UDARC problem, we provide two domain adaptation models. The first one learns the out-domain LM and in-domain RC task sequentially. The second one is the proposed model that uses a multi-task learning approach of LM and RC. The models can retain both the RC capability acquired from the supervised data in the source domain and the LM capability from the unlabeled data in the target domain. We evaluated the models on UDARC with five datasets in different domains. The models outperformed the model without domain adaptation. In particular, the proposed model yielded an improvement of 4.3/4.2 points in EM/F1 in an unseen biomedical domain.

Keywords: Reading Comprehension, Domain Adaptation, Unsupervised Learning

1. Introduction

Reading comprehension (RC) is a task to acquire a capability of understanding natural language for question answering with textual sources. It has seen significant progress since the release of numerous datasets such as SQuAD (Rajpurkar et al., 2016) and the rise of the deep neural models such as BiDAF (Seo et al., 2017). Recently, fine-tuning of pre-trained language models (LM) such as BERT (Devlin et al., 2019) has achieved state-of-the-art performance in many NLP tasks including RC.

However, such state-of-the-art models still do not have general linguistic intelligence; e.g., their accuracy is sensitively affected by the difference in the distribution between the training and evaluation datasets, such as in the domains of textual sources. This discrepancy becomes an issue in a real-world scenario. For instance, a business intending to introduce an RC application for a service must create tens of thousands of (passage, query, answer) tuples per domain of the service (Yogatama et al., 2019). However, such a large annotation to create training data costs much money and takes up a lot of time. When the domain entails personal information or expert knowledge, even crowd-sourcing can not be used to create the training data.

This study tackles a task, called **Unsupervised Domain Adaptation of Reading Comprehension** (UDARC). Table 1 shows the task setting of UDARC. The model can use the (passage, query, answer) tuples of the source domain for training. In addition, the model can use unlabeled passages for training in the target domain. Annotated corpora with QA pairs in the target domain can not be used for training. In addition to the interest in the domain discrepancy, we think that this setting is a natural one in real-world scenarios, where it is assumed that the RC application provider has documents that can be used as knowledge sources for RC in the target domain (e.g., technical documents about a product). UDARC thus enables people who have such doc-

	Training			Evaluation		
	Input Passage	Query	Output Answer	Input Passage	Query	Output Answer
Source	✓	✓	✓			
Target	✓			✓	✓	✓

Table 1: Task setting of unsupervised domain adaptation for reading comprehension (UDARC). Only unlabeled passages can be used for training in the target domain. In terms of QA pairs, the task requires zero-shot domain adaptation.

uments in the target domain but have no RC training data for the domain to introduce RC applications. This scenario also applies to low-resource languages, where the availability of RC training data is limited.

We hypothesize that the poor performance of in-domain models for out-domains is caused by a lack of LM capability for out-domains. That is, we can improve the question answering accuracy, without the RC data for the out-domains, by training the language model with textual sources about the out-domains in an unsupervised fashion. In this study, we introduce a no-adaptation baseline model. It transfers a BERT model fine-tuned with the source domain RC dataset to the target domain without domain adaptation. As a natural unsupervised domain adaptation approach, we investigate a sequential model. It adapts the language model of BERT with the unlabeled passages in the target domain, and then it fine-tunes BERT with the source domain RC dataset.

Moreover, we propose the multi-task learning approach of LM in the target domain and RC in the source domain. It is more promising because the model avoids forgetting about the target domain while fine-tuning. This study investigates the feasibility of UDARC and the effectiveness of these models on various domain datasets. Our main con-

tributions are as follows.

- We tackled unsupervised domain adaptation of reading comprehension (UDARC). To solve UDARC, we hypothesize that the ability to understand the out-domain passages can be learned without supervised annotation for RC.
- We propose a multi-task learning approach of the LM in the target domain and the RC in the source domain. This approach retains both knowledge of RC in the source domain and knowledge of LM in the target domain without forgetting. We use BERT as the backbone model.
- We evaluated three models on UDARC with five datasets in different domains. The domain adaptation models outperformed the no-adaptation model. In particular, the proposed model (adapted from the Wikipedia domain to a biomedical domain) yielded the best performance, with a 4.3/4.2 point gain in EM/F1 over the model without domain adaptation.
- We thoroughly investigated cases in which UDARC is effective. We found that the task is promising when the target domain is unseen in the training of the source domain and the pre-training of the language model and when the training data in the target domain are insufficient to acquire the RC and LM capabilities.

2. Problem Formulation

Here, we will focus on unsupervised domain adaptation of extractive RC, which is the most popular problem formulation of RC. The UDARC task is defined as follows.

PROBLEM 1. Let I_S be instances of 3-tuples (passage, query, answer span) in the *source* domain, and I_T be unlabeled passages in the *target* domain. The task of UDARC is to accurately answer the query (extract the correct answer-span from an input passage) in the target domain by training an extractive RC model with I_S and I_T .

3. Related Work

3.1. Reading Comprehension

Golub et al. (2017) and Wang et al. (2019) also tackled UDARC. Their approaches are to create pseudo QA pairs for training in the target domain. They use the answer extraction model to find a potential answer from the passage and the query generation model to create the query given the potential answer and the passage. Our approach is different from theirs because, according to our hypothesis, we can acquire knowledge in the target domain without RC training. An advantage of our approach is the low computational cost due to the lack of a need for the training of the Sequence-to-Sequence model to generate the query.

The MRQA 2019 shared task has a similar motivation as ours (Fisch et al., 2019). Its goal is to generalize to new test distributions and be robust to test-time perturbations. The task is on six in-domain training data and twelve out-domain evaluation data (a many-to-many setting). Our motivation is to acquire the ability to understand out-domain

passages from unlabeled passages and the ability to answer a query from annotated RC training data for domain adaptation (a one-to-one setting). This motivation corresponds to a real-world scenario. The shared task resulted in the certification of the dependence of RC performance on the backbone LM capability. This result supports the effectiveness of our approach to improve LM for domain adaptation. RC is essentially classified into four types: extractive (Rajpurkar et al., 2016), multiple-choice (Lai et al., 2017), generative (Nguyen et al., 2016), and cloze-style (Hermann et al., 2015). Although this study focused on extractive RC, we believe that the concept of our model can be applied to other types of RC.

Many RC datasets have been published recently, but the ones for closed domains are limited in number. There are datasets, for example, in biomedical domain (Šuster and Daelemans, 2018), scientific domain (Clark et al., 2018; Johannes Welbl and Gardner, 2017), and software domain (Dhingra et al., 2017). The limited number is one reason for our developing UDARC. Despite the demand for RC in closed domains, which requires expert knowledge, the annotation cost of the dataset makes it difficult to create the dataset.

3.2. Domain Adaptation

Unsupervised domain adaptation is a task to adapt the model to the target domain with labeled source data and unlabeled target data. While supervised domain adaptation with labeled target data first trains the model in the source domain and then adapts it to the target domain, unsupervised domain adaptation takes another approach. In NLP, Ziser and Reichart (2019) uses a two-step algorithm; the model first learns the representations and then learns the classification. Miller (2019) uses the multi-task learning approach of the classification and the feature selection of structural correspondence learning (Blitzer et al., 2006). Here, Ganin et al. (2016) trains the representation learning and the classification jointly with a domain-adversarial neural network.

However, UDARC can not use standard unsupervised domain adaptation methods. As shown in Table 1, an instance is divided into an input (e.g., a (passage, query) pair in RC) and output (e.g., an answer in RC). Unsupervised domain adaptation allows inputs to be used in the target domain without outputs. In RC, neither the output nor the queries in the inputs can be used for training.

Zero-shot learning is a task to adapt the model so that it can predict unseen classes in the training (Socher et al., 2013). In particular, the recently proposed zero-shot domain adaptation (Yang and Hospedales, 2015; Peng et al., 2018) can not use the inputs for the training of domain adaptation. UDARC can be interpreted as a kind of zero-shot domain adaptation, because it can not use queries in the inputs for the training.

Zero-shot domain adaptation is a more challenging task than unsupervised domain adaptation, and there are few studies on it. Peng et al. (2018) allows task-irrelevant data in the target domain. This is similar to our setting, but their task is classification of images, which is rather different from extractive RC, which is a special

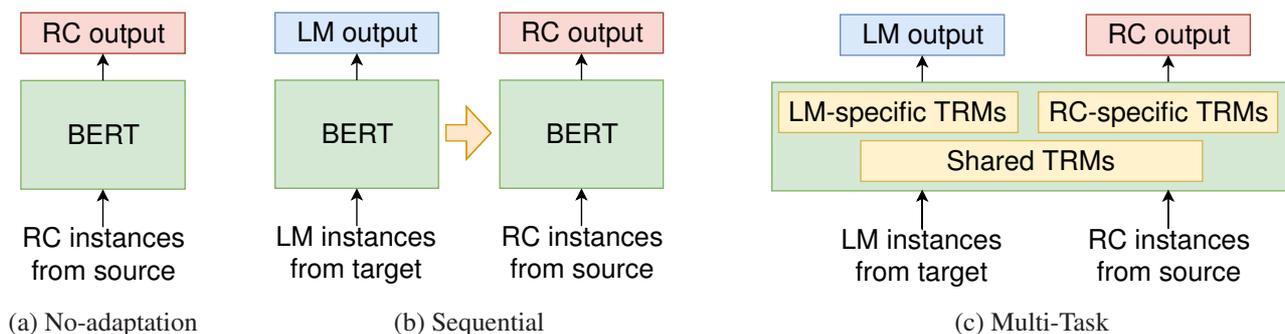


Figure 1: Overview of the three models. The LM instances are from the target domain, and the RC instances are from the source domain. “TRMs” means transformer layers.

case of sequence labeling for text. The results of other studies have limitations when they are used for UDARC. Yang and Hospedales (2015) and Mancini et al. (2019) hypothesize that the training data are from multiple domains. Ishii et al. (2019) suppose there is prior knowledge about what factors cause the differences between the source and target data distributions.

4. Methods

This section introduces three models: the baseline model without domain adaptation, the sequential domain adaptation model, and the proposed model using multi-task learning. We used a pre-trained BERT_{base} model as the backbone architecture for the three models. It was trained with large-scale corpora: BookCorpus (800M words) (Zhu et al., 2015) and English Wikipedia (2,500M words). It has achieved state-of-the-art performance on various RC datasets. See (Devlin et al., 2019) for the details of BERT_{base}. Note that the models can use other pre-trained LM consisting of stacked layers as their backbone architecture.

4.1. No-adaptation Baseline

This baseline model is a BERT fine-tuned with an RC dataset in the source domain. It does not consider domain adaptation, so its performance corresponds to the lower bound of UDARC.

Figure 1 (a) shows an overview of the model. We follow the fine-tuning setup for extractive RC, as in (Devlin et al., 2019). We add a linear layer for extractive RC on top of the BERT layers, where its output dimension is two. The first dimension represents a score that the token is the start of the answer span, and the other dimension represents the end. The input sequence is [‘[CLS]’; *query*; ‘[SEP]’; *passage*; ‘[SEP]’], where ‘[CLS]’ and ‘[SEP]’ are special tokens and ‘;’ means concatenation.

4.2. Sequential Model

This model first adapts the pre-trained BERT with unsupervised passages in the target domain. Then, it fine-tunes the adapted BERT with the RC dataset of the source domain. Figure 1 (b) shows an overview of the model. For the first unsupervised adaptation, a linear layer for LM is added at top of the BERT layers. The BERT including the LM

output layer is trained (with the sentences in the unsupervised passages) in the same manner as in the pre-training of BERT, i.e., by using masked language modeling (MLM) and next sentence prediction (NSP). The input sequence is [‘[CLS]’; *sentences 1*; ‘[SEP]’; *sentences 2*; ‘[SEP]’]. After the adaptation, another linear layer for RC is added, and the BERT including the RC output layer is fine-tuned in the same manner as in the no-adaptation baseline.

Note that we may conduct domain adaptation in reverse order, i.e., first build an RC model in the source domain and then adapt it to the target domain. However, this order causes catastrophic forgetting of the capability of finding an answer span. The sequential model follows the pipeline strategy of the previous unsupervised domain adaptation work (Ziser and Reichart, 2019). Their model learns representations first and classification after that, but their task is not RC and they do not use a pre-trained language model.

4.3. Proposed Multi-Task Model

We consider that the sequential model forgets the knowledge in the target domain while it is being fine-tuned with RC training data. Moreover, it is known that multi-task fine-tuning of the pre-trained language models outperforms single-task fine-tuning (Liu et al., 2019a). For this reason, we propose a multi-task learning approach. Multi-task learning adds an RC linear layer and LM linear layer to the top of the BERT layers. The model uses the RC linear layer for RC instances and the LM linear layer for LM instances as the output layer. Figure 1 (c) shows an overview of the model.

Our multi-task learning approach uses the following two techniques.

4.3.1. Using Shared and Specific Transformer Layers

First, we use the top- n Transformer layers (TRMs) for the tasks separately as task-specific RC or LM layers. The task-specific layers are cloned and initialized from the original pre-trained BERT_{base} layers. The other layers are commonly shared by the tasks. The RC instances pass through the shared layers, the RC-specific TRMs, and the RC output layer. The LM instances pass through the shared layers, the LM-specific TRMs, and the LM output layer. This idea follows Tenney et al. (2019)’s observation that basic syntactic information (e.g., part-of-speech tagging) is captured

Algorithm 1 Multi-task learning approach

Input: source RC instances I_S , target LM instances I_T , num. of steps N , the RC training ratio k

- 1: **for all** i in $1, \dots, N$ **do**
- 2: Select mini-batch $b_S \sim I_S$
- 3: Train the shared layers, RC-specific layers, and RC output layer with b_S .
- 4: **if** $i \% k == 0$ **then**
- 5: Select mini-batch $b_T \sim I_T$
- 6: Train the shared layers, LM-specific layers, and LM output layer with b_T .
- 7: **end if**
- 8: **end for**

in lower layers and high-level semantic information (e.g., coreference labeling) are captured in higher layers. We consider that the basic syntactic information is the common features between tasks, and the high-level semantic information is closely related to the output of the task. Therefore, the LM instances in the target domain should use the shared TRMs in order to capture the basic syntactic information in the target domain. The LM instances should not share the higher layers of BERT with the RC instances, because the way it uses the high-level semantic information is different from in RC.

4.3.2. One-Segment LM Training

Second, we preprocess each LM instance as a sequence with one segment. That is, each LM instance is $['[CLS]'; '[LM]'; passage; '[SEP]']$, and the segment ids are a zero vector. Therefore, NSP is not used for training of our multi-task model. This intends to prevent the learning of the segment interaction in NSP from disturbing the learning of query-passage interaction, because the two segments only interact in the RC training. $['[LM]']$ is a special token meaning that the instance is an LM instance.

4.3.3. Training Procedure

We perform the RC training and LM training alternately. The RC training is performed k times as many times as the LM training. The algorithm is shown in Algorithm 1. Note that the model size and computational time in the evaluation are the same as in the original BERT, because the LM-specific layers are not used for evaluation. In the training, the computational time remains about the same, because the training for each mini-batch is the same as in the original BERT.

5. Experiments

5.1. Dataset

We evaluated the UDARC task on various datasets. The training data in the source domain should be large and cover a wide range of topics. Moreover, the test data in the target domain should be from a closed domain. Here, we selected five datasets from different domains. Table 2 shows the statistics. We used the development data for the evaluation because some of the datasets do not include test data. Note that not all the unlabeled passages were used. The number of used unlabeled passages depended on the experimental

dataset	domain	# training data	# dev. data	# unlabeled passages
SQuAD	Wikipedia	87599	10570	19047
NewsQA	news	107064	5988	95933
BioASQ	biomedical	0	1504	55148
DuoRC	movie	69524	15591	5137
Natural Questions	HTML Wikipedia	104071	12836	12222

Table 2: Statistics of the datasets used in the experiments.

setup (the number of fine-tuning epochs and the value of k) and the size of the training dataset.

SQuAD1.1 is an RC dataset from Wikipedia (Rajpurkar et al., 2016). We used this dataset as the source domain or target domain. We used the passages of the training data as the unlabeled passages when the dataset was in the target domain.

NewsQA is an RC dataset from CNN news (Trischler et al., 2017). We used this dataset as the source domain or target domain. We used the CNN news scripts (Hermann et al., 2015) as unlabeled passages similar to the data collection of NewsQA.

BioASQ is a biomedical semantic indexing and question answering challenge¹ (Tsatsaronis et al., 2015). The MRQA 2019 shared task preprocesses this dataset for extractive RC; we used the MRQA version of this dataset. The training data of this dataset is not provided in the MRQA 2019 shared task. We used this dataset only for the target domain. We collected the unlabeled passages from the abstracts of PubMed articles.

Our main interest in the experiments is the performance of the unsupervised domain adaptation models on BioASQ. This is because the BioASQ domain is not fully covered by the source domain (Wikipedia or news) or by the BERT pre-training (BookCorpus and Wikipedia).

DuoRC is an RC dataset in the movie domain (Saha et al., 2018). DuoRC provides parallel movie plots from Wikipedia and IMDb. In the experiment, we used the ParaphraseRC task in DuoRC, where each query is created on a Wikipedia movie article, and each passage is collected from an IMDb article corresponding to the same movie. The ParaphraseRC task is difficult because it is designed to contain a large number of queries with low lexical overlap between queries and their corresponding passages. We used this dataset only for the target domain and the passages of the training data as the unlabeled passages.

The movie domain is different from the source domains, whereas the BERT pre-training covers many stories in BookCorpus. Therefore, we think that the pre-trained BERT has acquired knowledge of language modeling for this dataset.

Natural Questions (NQ) is an RC dataset containing passages from Wikipedia written in HTML format, where

¹Task 7b, Biomedical Semantic QA, is held with ECML PKDD 2019. See <http://BioASQ.org/>.

	Train w./ SQuAD	Domain Adaptation	Target			
			NewsQA	BioASQ	DuoRC	NQ
Standard RC in Target			41.5/56.0	—	20.2/27.2	58.9/72.2
No-adaptation	✓		35.2/50.7	41.1/53.6	24.5/33.0	44.4/57.5
Sequential	✓	✓	35.2/51.0	44.5/57.1	25.4/33.8	—
Multi-Task	✓	✓	35.9/51.4	45.4/57.8	25.5/34.1	43.8/56.7

Table 3: Results when the SQuAD was the source dataset. EM is on the left and F1 is on the right in each cell. The top row is the supervised training in the target domain, so it is the expected upper bound of UDARC. The Standard RC in BioASQ is empty, due to the lack of training data. The sequential model for NQ is empty. NSP of the BERT pre-training can not be applied to HTML with a lot of HTML tags (e.g., List and Table).

	Train w./ NewsQA	Domain Adaptation	Target			
			SQuAD	BioASQ	DuoRC	NQ
Standard RC in Target			80.9/88.4	—	20.2/27.2	58.9/72.2
No-adaptation	✓		59.8/73.9	34.5/48.3	22.5/31.2	39.0/52.7
Sequential	✓	✓	59.7/75.3	36.6/ 50.4	23.7/ 32.7	—
Multi-Task	✓	✓	60.6/75.8	36.8/50.3	23.8/32.3	42.0/56.2

Table 4: Results when the NewsQA was the source dataset.

each passage is given as a sequence of words and HTML tags (Kwiatkowski et al., 2019). We used the preprocessed dataset for extractive RC in the MRQA 2019 shared task for the training and evaluation. We used the passages of the training data in the original NQ as the unlabeled passages. Although the domain is the same as that of SQuAD and is also covered by the BERT pre-training, we used this dataset to confirm whether our model can adapt to HTML format without supervised RC data on the HTML format.

5.2. Experimental Setup

We compared three models, no-adaptation, sequential, and multi-task on the above datasets.

We used the PyTorch implementation of BERT². We trained the models on four NVIDIA Tesla P100 GPUs. The optimizer was Adam (Kingma and Ba, 2014). The warm-up proportion was 0.1 and the learning rate was 0.00005. The batch size was 32. There were three epochs. The input length was 384. Sequences longer than the input length were truncated with a stride length of 128. The other hyperparameters followed those of BERT_{base}. The ratio of RC training to LM training k is 10. The number of task-specific layers n is 3. The hyperparameters are fixed in all of the fine-tunings for RC. The hyperparameters of the pre-training in the sequential model follow the default settings of the implementation, except that the input length is 512, which is the longest case.

We evaluated the answer prediction in terms of exact match (EM) and partial match (F1). These are official metrics of SQuAD.

5.3. Results

Under what condition is UDARC effective? Table 3 (the source domain is Wikipedia from SQuAD) and Table 4 (news from NewsQA) show the performance of the models for the target domain. The ‘‘Standard RC in Target’’ row

lists the results in the standard RC (non-UDARC) setting, where each model was trained and evaluated in the target domain, so these are expected to be the upper bounds. First, we discuss each target-only dataset separately.

BioASQ. BioASQ is of main interest in our experiments. The two domain adaptation models outperformed the no-adaptation baseline in both the source domain settings. The proposed model improved on the no-adaptation model trained in SQuAD by 4.3/4.2 points.

BioASQ is in a domain that is not included in the BERT pre-training corpora. The results showed that the UDARC framework is effective in adapting to unseen domains. Moreover, this result confirms that our hypothesis behind UDARC is correct; i.e., the ability to understand the out-domain passages can be learned from a non-annotated corpus and the ability to answer the query can be learned from annotated RC training data, even though BERT is pre-trained with very large corpora.

DuoRC. All unsupervised domain adaptation models outperformed the model trained with the supervised dataset in the target domain³. This surprising result is caused by the difficulty of the ParaphraseRC task of DuoRC. It has low lexical overlap between queries and their corresponding passages, and the passages are rather long and complicated. As a result, the training data are too difficult for learning the RC capability. We think UDARC is promising when it is difficult to acquire the LM and RC capabilities with the supervised datasets.

Natural Questions. We compared four models (No-adaptation / Multi-Task with the source domain of SQuAD / NewsQA) in the target domain of NQ. The results indicated that unsupervised adaptation from SQuAD to NQ did not

³The performance of the supervised model was 20.2/27.2, which is similar to the performance (19.7/27.6) reported in the original paper.

²<https://github.com/huggingface/pytorch-transformers>

	NewsQA	BioASQ	DuoRC	NQ
Standard RC	80.9/88.4			
Sequential	81.2/88.6	80.6/88.4	81.0/88.4	—
Multi-Task	81.1/88.5	81.1/88.5	80.7/88.3	80.9/88.4

Table 5: Results for when the source and evaluation datasets were SQuAD. The model was adapted to each target dataset. The top row is the no-adaptation model trained only with the SQuAD dataset.

	SQuAD	BioASQ	DuoRC	NQ
Standard RC	41.5/56.0			
Sequential	41.8/56.9	42.0/57.1	42.7/58.0	—
Multi-Task	42.6/57.6	42.0/57.0	42.8/57.8	42.3/57.4

Table 6: Results for when the source and the evaluation datasets were NewsQA. The model was adapted to each target dataset. The top row is the no-adaptation model trained only with the NewsQA dataset.

improve accuracy; on the other hand, unsupervised adaptation from NewsQA to NQ was effective. These results can be interpreted as meaning that the multi-task approach trained in the news domain as the source successfully adapted to the Wikipedia domain. However, the proposed model trained with plain text failed to adapt to HTML format. We consider that the adaptation to the HTML format is a more challenging task than domain adaptation. Here, the task design of the language modeling remains as future work to understand text in HTML format, such as understanding of the dependencies among segments separated by HTML tags.

What is the performance of the three models? Here, let us discuss the model performances shown in Table 3 and Table 4. Except for NQ, the domain adaptation models outperformed the no-adaptation baseline. In terms of the EM metric, the multi-task model outperformed the sequential model for all source/target settings. This tendency showed the possibility that the sequential model forgets the out-domain knowledge while it is being fine-tuned. However, there was no statistical significance between the two models. We observed that the sequential model was as effective as the multi-task model. It is worth pre-training BERT with the unlabeled target domain data after pre-training in the general domain. This finding can be applied to other NLP tasks.

In comparison to the related work, Wang et al. (2019) only refers to the experiments with BERT in the setting from SQuAD to NewsQA. The improvement of their domain adaptation method over BERT fine-tuned in the source domain is 0.6/0.5 points, which is comparable to our 0.7/0.7 point improvement, though their pseudo QA generation approach requires more computational cost.

Does domain adaption hurt the performance in the source domain? We evaluated the drop in performance in the source domain due to the domain adaptation. Table

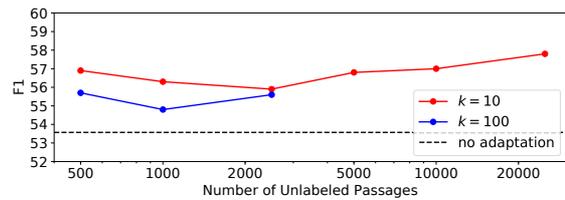


Figure 2: Performance (F1) of the proposed model in adaptation from SQuAD to BioASQ in terms of the number of unlabeled passages. The proposed model was trained with k times fewer LM instances than RC instances. The right-most coordinate of each line is the maximum number of used LM instances determined by k , the number of epochs, and the number of training data.

5 and Table 6 show the results. The “Standard RC” row is the model trained only in the source domain.

Surprisingly, the sequential model and the multi-task model tend to outperform the standard RC training in the source domain. In addition to the discrepancy between domains, a discrepancy also exists between the training samples and the evaluation samples. We consider that the domain adaptation has the effect of generalization, so the model overcomes the sampling discrepancy. UDARC requires no additional supervised data for training, so the framework of UDARC can be easily extended to the standard RC task in which we can expect an improvement in performance.

How many unlabeled passages are required? We evaluated the performance of the proposed model in terms of the number of unlabeled passages. The experiments were on adaptation from SQuAD to BioASQ, because this setting showed the largest improvements and BioASQ is the main focus of UDARC. Figure 2 shows the results.

The results show that the proposed model outperformed the no-adaptation baseline even when there were only 500 unlabeled passages. The gain was 3.7/3.3 points in EM/F1. In terms of the RC-LM ratio k , $k = 10$ was preferred, but we should note that in some cases, $k = 100$ was preferred in the pilot experiments. The biomedical domain is far from the pre-training corpus of BERT, so we consider that the moderate frequency of LM training steps is larger than in other domains included in the corpora. We found that the performance does not always increase as more unlabeled passages come to be used, though the best performance is with the full 26280 passages.

What domain is preferred as the source domain? To evaluate the preference about the source domain under the same conditions, we equalized the number of training data in the source domain. Figure 3 and Figure 4 shows the performance of the proposed model with BioASQ and DuoRC as the target dataset.

On BioASQ, the proposed model adapted from SQuAD outperformed the model adapted from NewsQA. In contrast, on DuoRC, the performance was on par. Therefore, the performance of the proposed model depends on the selection of the source domain, but the preferred source domain cannot as yet be identified.

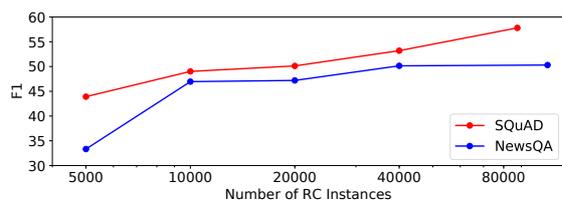


Figure 3: Performance (F1) of the proposed model in adaptation from SQuAD and NewsQA to BioASQ versus number of source training data.

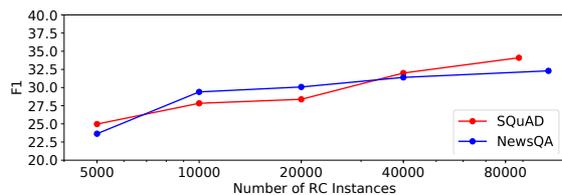


Figure 4: Performance (F1) of the proposed model in adaptation from SQuAD and NewsQA to DuoRC versus number of source training data.

How much source data are required? Figure 3 and Figure 4 show the performance in terms of the number of source training data. The results indicate that more training data results in higher performance. In particular, the improvement grows rapidly until 10000 instances, after which it becomes slower. This result coincides with the observation of Yogatama et al. (2019). They showed that tens of thousands of training data are required to fine-tune BERT. We consider that the required number of source data for UDARC shows the same tendency as in the standard RC task.

6. Conclusion and Future Work

This paper studied UDARC to adapt an RC model to the target domain without any annotated data in the target domain.

This is the first study to focus on the unsupervised domain adaptation for acquiring the ability to answer the question from the RC task in the source domain and the ability to understand the out-domain passages from the LM task in the target domain. This approach is different from the related work, which generates the pseudo QA pairs for the training. We described two unsupervised domain adaptation models using BERT. In addition, the proposed model reduces the forgetting of out-domain knowledge while it is being fine-tuned.

We evaluated the two models and the no-adaptation model on the five datasets in different domains. As a result, the domain adaptation models outperformed the no-adaptation model especially well when the target domain was not contained in the source domain or the BERT pre-training corpora. The proposed model (adapted from the Wikipedia domain to the biomedical domain) yielded the best performance, with a 4.3/4.2 points gain in EM/F1 over the no-adaptation model. We believe that this study sheds light on

the importance and feasibility of UDARC. Our experiments also showed that the UDARC framework has the potential to outperform a model trained with supervised datasets in the target domain when it is difficult to acquire the LM and RC capabilities from the supervised datasets.

Pre-trained language models (Devlin et al., 2019; Radford et al., 2019; Yang et al., 2019; Liu et al., 2019b) and the fine-tuned models achieved state-of-the-art performance in many NLP tasks, including RC. We believe that this study considering unsupervised domain adaptation with BERT will foster great contributions to various fields of NLP that have not been the subject of previous work in unsupervised domain adaptation.

7. Bibliographical References

- Blitzer, J., McDonald, R., and Pereira, F. (2006). Domain adaptation with structural correspondence learning. In *EMNLP*, pages 120–128.
- Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. (2018). Think you have solved question answering? Try ARC, the AI2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186.
- Dhingra, B., Mazaitis, K., and Cohen, W. W. (2017). Quasar: Datasets for question answering by search and reading. *arXiv preprint arXiv:1707.03904*.
- Fisch, A., Talmor, A., Seo, M., Choi, E., and Chen, D. (2019). MRQA 2019 shared task: Evaluating generalization in reading comprehension. In *MRQA@EMNLP*, pages 1–13.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. (2016). Domain-adversarial training of neural networks. *JMLR*, 17(1):2096–2030.
- Golub, D., Huang, P.-S., He, X., and Deng, L. (2017). Two-stage synthesis networks for transfer learning in machine comprehension. In *EMNLP*, pages 835–844.
- Hermann, K. M., Kocisky, T., Grefenstette, E., Espeholt, L., Kay, W., Suleyman, M., and Blunsom, P. (2015). Teaching machines to read and comprehend. In *NIPS*, pages 1693–1701.
- Ishii, M., Takenouchi, T., and Sugiyama, M. (2019). Zero-shot domain adaptation based on attribute information. *arXiv preprint arXiv:1903.05312*.
- Johannes Welbl, N. F. L. and Gardner, M. (2017). Crowdsourcing multiple choice science questions. In *W-NUT@EMNLP*, pages 94–106.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., et al. (2019). Natural questions: a benchmark for question answering research. *TACL*, 7:453–466.
- Lai, G., Xie, Q., Liu, H., Yang, Y., and Hovy, E. (2017).

- RACE: Large-scale reading comprehension dataset from examinations. In *EMNLP*, pages 785–794.
- Liu, X., He, P., Chen, W., and Gao, J. (2019a). Multi-task deep neural networks for natural language understanding. In *ACL*, pages 4487–4496.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019b). RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Mancini, M., Bulò, S. R., Caputo, B., and Ricci, E. (2019). Adagraph: Unifying predictive and continuous domain adaptation through graphs. In *CVPR*, pages 6568–6577.
- Miller, T. (2019). Simplified neural unsupervised domain adaptation. In *NAACL-HLT*, pages 414–419.
- Nguyen, T., Rosenberg, M., Song, X., Gao, J., Tiwary, S., Majumder, R., and Deng, L. (2016). MS MARCO: A human generated machine reading comprehension dataset. In *CoCo@NIPS*.
- Peng, K.-C., Wu, Z., and Ernst, J. (2018). Zero-shot deep domain adaptation. In *ECCV*, pages 764–781.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. (2019). Language models are unsupervised multitask learners. *Technical Report OpenAI*.
- Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. (2016). SQuAD: 100,000+ questions for machine comprehension of text. In *EMNLP*, pages 2383–2392.
- Saha, A., Aralikkatte, R., Khapra, M. M., and Sankaranarayanan, K. (2018). DuoRC: Towards complex language understanding with paraphrased reading comprehension. In *ACL*, pages 1683–1693.
- Seo, M., Kembhavi, A., Farhadi, A., and Hajishirzi, H. (2017). Bidirectional attention flow for machine comprehension. In *ICLR*.
- Socher, R., Ganjoo, M., Manning, C. D., and Ng, A. (2013). Zero-shot learning through cross-modal transfer. In *NIPS*, pages 935–943.
- Šuster, S. and Daelemans, W. (2018). CliCR: A dataset of clinical case reports for machine reading comprehension. In *NAACL-HLT*, pages 1551–1563.
- Tenney, I., Das, D., and Pavlick, E. (2019). BERT rediscovered the classical nlp pipeline. In *ACL*, pages 4593–4601.
- Trischler, A., Wang, T., Yuan, X., Harris, J., Sordani, A., Bachman, P., and Suleman, K. (2017). NewsQA: A machine comprehension dataset. In *Rep4NLP@ACL*, pages 191–200.
- Tsatsaronis, G., Balikas, G., Malakasiotis, P., Partalas, I., Zschunke, M., Alvers, M. R., Weissenborn, D., Krithara, A., Petridis, S., Polychronopoulos, D., Almirantis, Y., Pavlopoulos, J., Baskiotis, N., Gallinari, P., Artieres, T., Ngonga, A., Heino, N., Gaussier, E., Barrio-Alvers, L., Schroeder, M., Androutsopoulos, I., and Paliouras, G. (2015). An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics*, 16:138.
- Wang, H., Gan, Z., Liu, X., Liu, J., Gao, J., and Wang, H. (2019). Adversarial domain adaptation for machine reading comprehension. In *EMNLP-IJCNLP*, pages 2510–2520.
- Yang, Y. and Hospedales, T. (2015). Zero-shot domain adaptation via kernel regression on the grassmannian. *DIFF-CV@BMVC*.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., and Le, Q. V. (2019). XLNet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.
- Yogatama, D., d’Aulume, C. d. M., Connor, J., Kocisky, T., Chrzanowski, M., Kong, L., Lazaridou, A., Ling, W., Yu, L., Dyer, C., and Blunsom, P. (2019). Learning and evaluating general linguistic intelligence. *arXiv preprint arXiv:1901.11373*.
- Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., and Fidler, S. (2015). Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *ICCV*, pages 19–27.
- Ziser, Y. and Reichart, R. (2019). Task refinement learning for improved accuracy and stability of unsupervised domain adaptation. In *ACL*, pages 5895–5906.