

Data and text mining

# BioADAPT-MRC: adversarial learning-based domain adaptation improves biomedical machine reading comprehension task

Maria Mahbub <sup>1,2,\*</sup>, Sudarshan Srinivasan<sup>2</sup>, Edmon Begoli<sup>2</sup>  
and Gregory D. Peterson<sup>1</sup>

<sup>1</sup>Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville, TN 37996, USA and <sup>2</sup>Cyber Resilience and Intelligence Division, Oak Ridge National Laboratory, Oak Ridge, TN 37830, USA

\*To whom correspondence should be addressed.

**Notice of Copyright:** This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

Associate Editor: Zhiyong Lu

Received on February 25, 2022; revised on May 31, 2022; editorial decision on July 2, 2022

## Abstract

**Motivation:** Biomedical machine reading comprehension (biomedical-MRC) aims to comprehend complex biomedical narratives and assist healthcare professionals in retrieving information from them. The high performance of modern neural network-based MRC systems depends on high-quality, large-scale, human-annotated training datasets. In the biomedical domain, a crucial challenge in creating such datasets is the requirement for domain knowledge, inducing the scarcity of labeled data and the need for transfer learning from the labeled general-purpose (source) domain to the biomedical (target) domain. However, there is a discrepancy in marginal distributions between the general-purpose and biomedical domains due to the variances in topics. Therefore, direct-transferring of learned representations from a model trained on a general-purpose domain to the biomedical domain can hurt the model's performance.

**Results:** We present an adversarial learning-based domain adaptation framework for the biomedical machine reading comprehension task (BioADAPT-MRC), a neural network-based method to address the discrepancies in the marginal distributions between the general and biomedical domain datasets. BioADAPT-MRC relaxes the need for generating pseudo labels for training a well-performing biomedical-MRC model. We extensively evaluate the performance of BioADAPT-MRC by comparing it with the best existing methods on three widely used benchmark biomedical-MRC datasets—BioASQ-7b, BioASQ-8b and BioASQ-9b. Our results suggest that without using any synthetic or human-annotated data from the biomedical domain, BioADAPT-MRC can achieve state-of-the-art performance on these datasets.

**Availability and implementation:** BioADAPT-MRC is freely available as an open-source project at <https://github.com/mmahbub/BioADAPT-MRC>.

**Contact:** mmahbub@vols.utk.edu

**Supplementary information:** [Supplementary data](#) are available at *Bioinformatics* online.

## 1 Introduction

During the consultation phase of primary patient care, for every two patients, healthcare professionals raise at least one question (Del Fiol *et al.*, 2014). Even though they can successfully find answers to 78% of the pursued questions, they never pursue half of their questions because of time constraints

and the suspicion that helpful answers do not exist, notwithstanding the availability of ample evidence (Bastian *et al.*, 2010; Del Fiol *et al.*, 2014). Additionally, searching existing resources for reliable, relevant and high-quality information poses an inconvenience for the clinicians on account of time limitation. This phenomenon elicits the dependency on general-information electronic

resources that are simple to use, such as Google (Hider et al., 2009). Apart from the healthcare professionals, there is also a growing public interest in learning about their medical conditions online (Fox and Duggan, 2013). Nevertheless, the criteria for ranking search results by general-purpose search engines does not conform directly to the fundamentals of evidence-based medicine and thus lacks rigor, reliability and quality (Hider et al., 2009).

While traditional information retrieval (IR) systems somewhat mitigate this issue, it still requires 4 h for a healthcare-information professional to find answers to queries related to complex biomedical resources (Russell-Rose and Chamberlain, 2017). Compared to the IR systems that usually provide the users (general population or healthcare professionals) a group of documents to interpret and find the exact answers, biomedical machine reading comprehension (biomedical-MRC) systems can provide exact answers to user inquiries, saving both time and effort.

Machine reading comprehension (MRC) is a challenging task in natural language processing (NLP), aiming to teach and evaluate the machines to understand user-defined questions, read and comprehend input contexts (namely, context) and return answers from them. The datasets in the MRC task consist of context–question–answer triplets where the question–answer pairs are considered labels. With the development and availability of efficient computing hardware resources, researchers have developed several state-of-the-art (SOTA) neural network-based (NN-based) MRC systems capable of achieving analogous or superior to human-level performance on several benchmark MRC datasets (Devlin et al., 2019; Joshi et al., 2017; Rajpurkar et al., 2016). However, this achievement is highly dependent on a large amount of high-quality human-annotated datasets that are used to train these systems (Rajpurkar et al., 2016). For domain-specific MRC tasks, especially biomedical-MRC, building a high-quality labeled dataset, specifically, the question–answer pairs residing in the dataset requires undeniable effort and knowledge of subject matter experts. This requirement leads to smaller biomedical-MRC datasets and, consequently, unreliably poor performance on the MRC task itself (Pergola et al., 2021). Hence, developing an approach that can effectively leverage unlabeled or small-scale labeled datasets in training the biomedical-MRC model is crucial for improving performance.

Researchers have addressed this issue by using transfer learning, a learning process to transfer knowledge from a source domain to a target domain (Pan and Yang, 2010). In domain-specific MRC problems, such as biomedical-MRC, the source domain is usually a general-purpose domain where a large-scale human-annotated MRC dataset is available. The target domain, in this case, is the biomedical domain. In this work, we focus on transferring the knowledge from an MRC model trained on a labeled general-purpose-domain dataset to the biomedical domain where only *unlabeled contexts* are available. Unlabeled contexts refer to only contexts in the MRC dataset with no question–answer pairs.

Often, directly transferring the knowledge representations (learned by an MRC model) from the source to the target domain can hurt the performance of the model because of the distributional discrepancies between the data seen at train and test time (Ganin and Lempitsky, 2015). Domain adaptation, a sub-setting of transfer learning (Pan and Yang, 2010), aims at mitigating these discrepancies through *simultaneous generation of feature representations that are discriminative from the viewpoint of the MRC task in the source domain and indiscriminative from the perspective of the shift in the marginal distributions between the source and target domains* (Ganin and Lempitsky, 2015).

We propose Adversarial learning-based Domain adAPTation framework for Biomedical Machine Reading Comprehension (BioADAPT-MRC), a new framework that uses adversarial learning to generate domain-invariant feature representations for better domain adaptation in biomedical-MRC models. In an adversarial learning framework, we train two adversaries (feature generator and discriminator) alternately or jointly against one another to generate domain-invariant features. Domain-invariant features imply that the feature representations extracted from the source- and the target-domain samples are closer in the embedding space.

While other recent domain adaptation approaches for the MRC task focus on generating pseudo question–answer pairs to augment the training data (Golub et al., 2017; Wang et al., 2019), we utilize only the unlabeled contexts from the target domain. This property makes our framework more suitable in cases where not only human-annotated dataset is scarce but also the generation of synthetic question–answer pairs is computationally expensive, and needs further validation from domain-experts (due to the sensitivity to the correctness of the domain knowledge).

We validate our proposed framework on three widely used benchmark datasets from the cornerstone challenge on biomedical question answering and semantic indexing, BioASQ (Tsatsaronis et al., 2015), using their recommended evaluation metrics. We empirically demonstrate that with the presence of no labeled data from the biomedical domain—synthetic or human-annotated—our framework can achieve SOTA performance on these datasets. We further evaluate the domain adaptation capability of our framework by using clustering and dimensionality reduction techniques. Additionally, we extend our framework to a semi-supervised setting and use varying ratios of labeled target-domain data for evaluation. Last but not least, we perform a thorough error analysis of our proposed framework to demonstrate its strengths and weaknesses.

The primary contributions of the article are as follows: (i) we propose BioADAPT-MRC, an adversarial learning-based domain adaptation framework that incorporates a domain-similarity discriminator with an auxiliary task layer and aims at reducing the domain shift between high-resource general-purpose domain and low-resource biomedical domain. (ii) We leverage the unlabeled contexts from the biomedical domain and thus relax the need for synthetic or human-annotated labels for target-domain data. (iii) We further extend the learning paradigm of BioADAPT-MRC to a semi-supervised setting. We show that our framework can be successfully employed to improve the performance of a pre-trained language model (PLM) in the presence of varying ratios of labeled target-domain data. (iv) Through comprehensive evaluations and analyses on several benchmark datasets, we demonstrate the effectiveness of our proposed framework and its domain adaptation capability for biomedical-MRC.

## 2 Background and related work

In this article, we focus on the biomedical-MRC task using the adversarial learning-based domain adaptation technique. Thus, our work is in the confluence of two main research areas: biomedical-MRC and domain adaptation using adversarial learning.

**Biomedical-MRC.** In the biomedical-MRC task, the goal is to extract an answer span, given a user-defined question and a biomedical context. In NN-based biomedical-MRC systems, the question–context pairs are converted from discrete textual form to continuous high-dimensional vector form using word-embedding algorithms, such as word2vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014), FastText (Bojanowski et al., 2017) and Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019), etc. Among numerous architectural varieties of these NN-based MRC systems, the transformer-based PLMs, such as BERT, are the current SOTA (Gu et al., 2021). The original BERT model is pre-trained on general-purpose English corpora. Considering the semantic and syntactic uniqueness of the biomedical text, researchers have developed different variants of BERT models for the biomedical domain that are pre-trained on several biomedical corpora, such as PMC full articles, PubMed abstracts and MMIC datasets. Some examples of such PLMs are BioBERT (Lee et al., 2020), PubMedBERT (Gu et al., 2021) and BioElectra (Raj Kanakarajan et al., 2021), which reportedly outperform the original BERT model in various biomedical NLP tasks. These PLMs are used as trainable encoding modules (encoders) in downstream biomedical NLP tasks, such as biomedical named entity recognition (NER) (Naseem et al., 2021), clinical-note classification (Agnikula Kshatriya et al., 2021), MRC (Jeong et al., 2020), etc. Usually, to accomplish the downstream tasks, such as biomedical-MRC by transferring the knowledge from the PLMs, researchers add a few task-specific layers, commonly feed-forward neural network layers, at the end of the encoders (Hosein et al., 2019; Jeong et al., 2020; Lee et al., 2020).

**Transfer learning.** Transfer learning is an approach to transfer knowledge representations acquired from a widely explored domain/task (source), to a new or less explored domain/task (target) (Pan and Yang, 2010). Adopting the notations provided by Pan and Yang (2010), a domain  $\mathcal{D}$  consists of a feature space  $\mathcal{X}$  (different from the feature representation learned by the network) and a marginal distribution of the learning samples  $X \in \mathcal{X}$ ,  $p(X)$ . In NLP, the marginal distributions are different when the languages are the same but the topics are different in the source and target domains (Bashath et al., 2022). Considering a label space  $\mathcal{Y}$ , for a given domain  $\mathcal{D}$ , a task  $\mathcal{T}$  can be described as  $(\mathcal{Y}, \mathcal{F}(\cdot))$ , where  $\mathcal{F}(\cdot)$  is a predictor function learned from the training data  $(X \in \mathcal{X}, Y \in \mathcal{Y})$ .

Data scarcity in the target domain is often an impediment to the model's performance while training for a target task. Transfer learning tackles this

issue by aiming to improve generalizability on a target task using acquired knowledge from a source domain  $\mathcal{D}_s$ , a target domain  $\mathcal{D}_t$  along with their respective associated tasks  $\mathcal{T}_s$  and  $\mathcal{T}_t$ . For this work, we use *transductive transfer learning*, which uses labeled source domain, unlabeled target domain, identical source and target tasks and different source and target domains. Depending on the similarity in the feature spaces, there are two cases of transductive transfer learning: (i) different feature spaces for the source and target domains, e.g. cross-lingual transfer learning and (ii) identical feature spaces, but different marginal probability distributions for the source- and target-domain samples, e.g. transfer learning between two domains with the same language but different topics (Bashath et al., 2022). In this article, we use the latter case of transductive transfer learning to train the biomedical-MRC system, otherwise known as domain adaptation (Pan and Yang, 2010).

**Domain adaptation.** Domain adaptation aims at increasing the generalizability of machine learning models when posed with unlabeled or very few labeled data from the target domain by generating domain-invariant representation (Glorot et al., 2011). One can enforce the learning of domain-invariant features in machine learning models by implementing the adversarial learning framework (Ganin and Lempitsky, 2015; Tzeng et al., 2017). In the adversarial setting, usually, a domain discriminator is incorporated into the MRC framework where besides performing the MRC task, the goal is to attempt at fooling the discriminator by generating domain-invariant features (Wang et al., 2019). Researchers have successfully applied domain adaptation in many tasks, such as sentiment classification (Glorot et al., 2011), speech recognition (Sun et al., 2017), neural machine translation (Thompson et al., 2019), NER (Vu et al., 2020) and image segmentation (Guan et al., 2021). However, compared to these tasks, the application of domain adaptation to the MRC task poses one additional challenge apart from missing answers—the missing questions in the target domain. Over recent years, researchers have proposed various methods to generate synthetic question-answer pairs from unlabeled contexts. For example, Wang et al. (2019) used NER and Bi-LSTM, Golub et al. (2017) used conditional probability, IOB tagger and Bi-LSTM and Yue et al. (2021) used seq2seq model with an attention mechanism to generate pseudo question-answer pairs. A multi-task learning approach has also been used for domain adaptation in MRC tasks (Nishida et al., 2020).

Among these works in MRC and domain adaptation, the AdaMRC model proposed by Wang et al. (2019) focuses on learning domain-invariant features in an adversarial setting as ours. However, the main differences between BioADAPT-MRC and AdaMRC are as follows: (i) while AdaMRC uses synthetically generated question-answer pairs to augment the target-domain dataset, BioADAPT-MRC directly uses the unlabeled contexts and thus relaxes the need for synthetic question-answer pairs. In later sections, we show that although synthetic questions can improve the performance of MRC tasks for various target domains, such as Wikipedia, web search log and news (Wang et al., 2019), they can hurt the performance of the MRC task for the biomedical domain. (ii) While AdaMRC uses the binary classification loss, BioADAPT-MRC uses triplet loss to minimize the domain shift between the source and target domains. Unlike binary classification loss, triplet loss considers both similarity and dissimilarity between two samples for gradient update and is known to be successful in deep metric learning where the aim is to map semantically similar instances closer in the embedding space and *vice versa* (Chen et al., 2018b; Kaya and Bilge, 2019; Kim et al., 2018). Moreover, triplet loss makes BioADAPT-MRC directly applicable to domain adaptation among more than two domains. While multiple prior works in computer vision have successfully used triplet loss for domain adaptation in numerous applications (Laiz et al., 2019; Wen et al., 2018), to the best of our knowledge, ours is the first in the application of the MRC task in NLP. (iii) To improve performance and stabilize the training process in the adversarial domain adaptation framework, BioADAPT-MRC uses an auxiliary task layer, similar to AC-GAN (Odena et al., 2017).

## 3 Materials and methods

In this section, we discuss our adversarial learning-based domain adaptation framework for the biomedical-MRC task.

### 3.1 Problem definition

Given an unlabeled target domain  $\mathcal{D}_t$  and a labeled source domain  $\mathcal{D}_s$  along with their respective learning tasks  $\mathcal{T}_t$  and  $\mathcal{T}_s$ , we assume that  $\mathcal{T}_t = \mathcal{T}_s$  and

$\mathcal{D}_t \neq \mathcal{D}_s$  because of  $p(X_t \in \mathcal{X}) \neq p(X_s \in \mathcal{X})$ , where  $p(\cdot)$  is the marginal probability distribution,  $X_t$  and  $X_s$  are learning samples from the target and source domains, respectively. Thus, while the tasks are identical, the domains are different due to different marginal probability distributions in their data.

In this work,  $\mathcal{D}_t$  is the biomedical domain where only unlabeled biomedical contexts are available, and  $\mathcal{D}_s$  is the general-purpose domain where large-scale labeled data are available. As mentioned in Section 2, despite having the same language, differences in the topics between two domains cause the domains to be different because of the dissimilarities in  $p(\cdot)$ . In this work, we consider that the general-purpose and the biomedical domains have different topics. Thus, we assume that  $p(X_t) \neq p(X_s)$ .

The task for both domains is extractive MRC. Given a question  $Q = \{q_1, q_2, \dots, q_n\}$  and a context  $C = \{c_1, c_2, \dots, c_m\}$ , extractive MRC predicts the start and end positions  $a^{\text{start}}$  and  $a^{\text{end}}$ , respectively, of the answer span  $A = \{c_i\}_{a^{\text{start}}}^{a^{\text{end}}}$  in  $C$  such that there exists one and only one answer span consisting of continuous tokens in the context. Here,  $q_i$  denotes the  $i$ th token in the question,  $c_i$  denotes the  $i$ th token in the context and  $n, m$ , respectively, denote the number of tokens in  $Q$  and  $C$ .

### 3.2 BioADAPT-MRC

Given the labeled and unlabeled inputs, respectively, from the source and target domains, our proposed framework BioADAPT-MRC aims at achieving the following two objectives: (i) predicting the answer spans from the provided contexts and (ii) addressing the discrepancies in the marginal distributions between the inputs in the source and target domains by generating domain-invariant features. Figure 1 demonstrates the three primary components of the BioADAPT-MRC framework:

- *Feature extractor* accepts a text sequence and encodes it into a high-dimensional continuous vector representation.
- *MRC-module* accepts the encoded representation from either the source domain (training time) or the target domain (test time), then predicts the start and end positions of the answer span  $A$  in  $C$ .
- *Domain-similarity discriminator* accepts the encoded representations from the source and target domains and learns to distinguish between them.

#### 3.2.1 Feature extractor

Given an input sample from either domain, the feature extractor  $M_F(\cdot)$  maps it to a common feature space  $F$ :

$$f_i = M_F(X_i), X_i \in \mathcal{D}_s \cup \mathcal{D}_t. \quad (1)$$

Here,  $f_i$  is the extracted feature for the  $i$ th input sample  $X_i$  from either  $\mathcal{D}_s$  or  $\mathcal{D}_t$ . We utilize the encoder of the PLM BioELECTRA (Raj Kanakarajan et al., 2021) as the feature extractor. We choose BioELECTRA for the following reason: while biomedical domain-specific BERT models, such as BioBERT, SciBERT outperform the original BERT models in several biomedical NLP tasks (Alsentzer et al., 2019), BioELECTRA has the best performance scores on the *Biomedical Language Understanding and Reasoning Benchmark* (Gu et al., 2021).

As mentioned in Section 2, the features in this task are trainable, high-dimensional word embeddings extracted from the question-context pairs. To generate these word embeddings, the BioELECTRA model utilizes the transformer-based architecture from one of the BERT-variants, ELECTRA. The ELECTRA model has 12 layers, 768 hidden size, 3072 feed-forward network (FFN) inner hidden size and 12 attention heads per layer (Clark et al., 2020). The pre-training corpora for BioELECTRA are 3.2 million PubMed Central full-text articles and 22 million PubMed abstracts, and the pre-training task is the replaced token prediction task. BioELECTRA has a vocabulary of size 30 522. The maximum number of tokens per input can be 512, where the embedding dimension of each token is 768. For each pair, the final tokenized input of the BioELECTRA model is  $\{[CLS], Q, [SEP], C, [SEP]\}$ . Here,  $Q$  and  $C$  are, respectively, the tokens from the question and the context,  $[CLS]$  is a special token that can be considered to have an accumulated representation of the input sequence (Devlin et al., 2019) and used for classification tasks,  $[SEP]$  is another special token that separates two consecutive sequences. Note that, since the samples in the target domain are unlabeled, in place of the question tokens  $Q$ , we use a

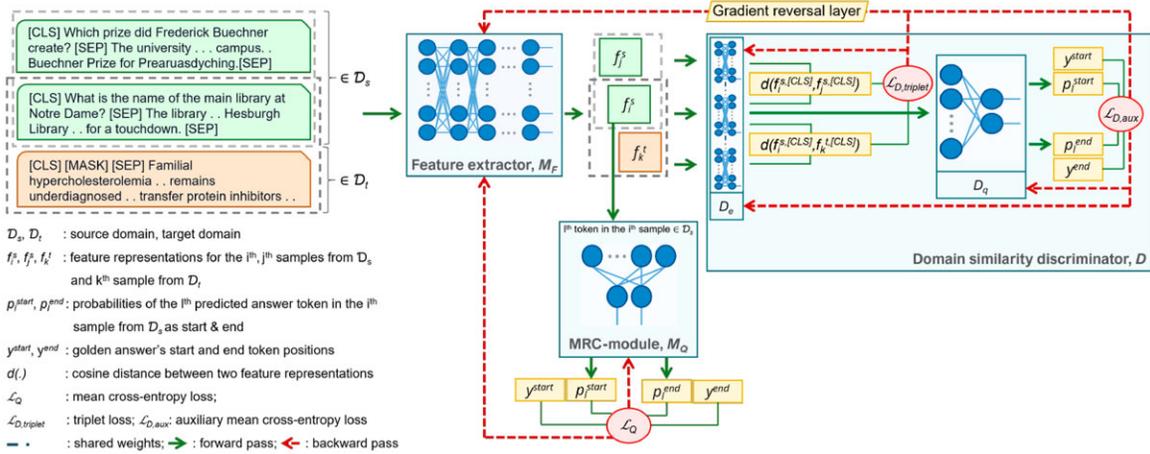


Fig. 1. BioADAPT-MRC: an BioADAPT-MRC task. The framework has three main components: (i) feature extractor  $M_F$ , (ii) MRC-module  $M_Q$  and (iii) domain-similarity discriminator  $D$

special token [MASK] to maintain consistency in the structure of the tokenized samples.

### 3.2.2 MRC-module

As the MRC-module  $M_Q(\cdot)$ , we add a simple fully connected layer with hidden size  $H = 768$  on top of the feature extractor  $M_F(\cdot)$  and use the softmax activation function to generate probability distributions for start and end token positions following Equation (2).

$$p_l^{start} = \frac{\exp(W^{start} \cdot h_l)}{\sum_{k=1}^{n_{seq}} \exp(W^{start} \cdot h_k)}; \quad p_l^{end} = \frac{\exp(W^{end} \cdot h_l)}{\sum_{k=1}^{n_{seq}} \exp(W^{end} \cdot h_k)}. \quad (2)$$

Here,  $p_l^{start}$  and  $p_l^{end}$  are the probabilities of the  $l$ th token to be predicted as start and end, respectively,  $h_l \in \mathbb{R}^H$  is the hidden representation vector of the  $l$ th token,  $W^{start} \in \mathbb{R}^H$  and  $W^{end} \in \mathbb{R}^H$  are two trainable weight matrices,  $n_{seq}$  is the input sequence length. We use the cross-entropy loss  $\mathcal{L}_Q$  on the predicted answer positions as the objective function for the  $M_Q$ . Since for each answer span prediction, we get two predicted outputs for the start and end positions, we average the total cross-entropy loss as shown in Equation (3).

$$\mathcal{L}_Q = -\frac{1}{2} (\log p_{y^{start}}^{start} + \log p_{y^{end}}^{end}). \quad (3)$$

Here, the golden answer's start and end token positions are represented by  $y^{start}$  and  $y^{end}$ , respectively. During test phase, the predicted answer span is selected based on the positions of the highest probabilities from  $p_{k \in [1, n_{seq}]}^{start}$  and  $p_{k \in [1, n_{seq}]}^{end}$ .

### 3.2.3 Domain-similarity discriminator

The domain-similarity discriminator  $D(\cdot)$  addresses the domain variance between two domains (caused by the discrepancies in the marginal probability distributions), as follows: in the adversarial setting,  $D(\cdot)$  learns to distinguish between the feature representations of the source- and target-domain samples generated by the feature extractor.  $D(\cdot)$  then penalizes the feature extractor for producing domain-variant feature representations and thus promotes the generation of domain-invariant features.  $D(\cdot)$  uses cosine distance between the feature representations of the input samples to distinguish between the domains. We consider that two samples are closer in the embedding space and thus have a greater chance to be in the same domain if their feature representations have a smaller cosine distance between them and *vice versa*. The input of  $D(\cdot)$  is a triplet  $(f_k^s, f_i^s, f_j^s)$ , where  $f_k^s, f_i^s$  and  $f_j^s$  are, respectively, the feature representations of the  $k$ th sample  $X_k^s$  from the target domain and  $i$ th sample  $X_i^s$  and  $j$ th sample  $X_j^s$  from the source domain extracted by  $M_F(\cdot)$ . The triplet is then split into two distinct pairs, consisting of  $(f_i^s, f_k^s)$  and  $(f_j^s, f_i^s)$ . As indicated in Figure 1, upon receiving each triplet,  $D$  accomplishes two tasks: (i) measures the similarity between  $(f_i^s, f_k^s)$  and dissimilarity between  $(f_j^s, f_i^s)$  and (ii) performs MRC task similar to  $M_Q$  for the source sample  $X_i^s$ .

For the first task, we introduce a Siamese network (Bromley et al., 1993)  $D_e$  with a single transformer encoder layer.  $D_e$  acts as a function that helps estimate the similarity and dissimilarity between the received pairs. Considering the success of the BERT models in many NLP tasks, for the Siamese network, we adopt the same architecture as any encoder layer in the BERT<sub>Base</sub> model, which has 12 attention heads, 768 embedding dimensions, 3072 FFN inner hidden size with 10% dropout rate and 'GeLU' activation function. We encode the input pairs  $(f_i^s, f_k^s)$  and  $(f_j^s, f_i^s)$ , using the same encoder network  $D_e$ . Considering the role of the special token [CLS] as explained in Section 3.2.1, to let  $D(\cdot)$  differentiate whether the pairs are from the same domain or not, we extract the [CLS] token representations from  $D_e$  for  $f_k^s, f_i^s$ , and  $f_j^s$ . We then use these [CLS] token representations to calculate the domain similarity and dissimilarity via triplet loss function (Weinberger et al., 2005) and use it as the learning objective of the discriminator  $D(\cdot)$  as shown in Equation (4).

$$\mathcal{L}_{D,triplet} = \max\{d(f_i^{s,[CLS]}, f_j^{s,[CLS]}) \times d(f_i^{s,[CLS]}, f_k^{t,[CLS]}) + \alpha, 0\}. \quad (4)$$

Here,  $f_i^{s,[CLS]}, f_j^{s,[CLS]}, f_k^{t,[CLS]}$  are, respectively, the [CLS] token representations of  $i$ th and  $j$ th samples from the source domain and the  $k$ th sample from the target domain where  $i \neq j$ ,  $d(\cdot) = 1.0 - \text{CosSim}(\cdot)$  is the cosine distance where  $\text{CosSim}(\cdot)$  is the cosine similarity,  $\alpha$  is the non-negative margin representing the minimum difference between  $d(f_i^{s,[CLS]}, f_j^{s,[CLS]})$  and  $d(f_i^{s,[CLS]}, f_k^{t,[CLS]})$  that is required for the triplet loss to be zero. To optimize  $D(\cdot)$ ,  $\mathcal{L}_{D,triplet}$  aims at minimizing  $d(\cdot)$  between the samples from the source domain and maximizing  $d(\cdot)$  between the samples from the source and target domains. Using triplet loss, our discriminator efficiently employs both similar and dissimilar information extracted by the feature extractor component of the model.

While using triplet loss in the adversarial setting, we also consider that the triplet loss might make representations from the source-domain dissimilar. Therefore, as an additional experiment, we try optimizing the discriminator by minimizing a distance-based loss, which is equivalent to just minimizing the distance  $d(\cdot)$  between the source- and the target-domain sample representations in the adversarial setting. We demonstrate the comparison between these two approaches in Section 4.5.1.

Adversarial learning for domain adaptation is known to be unstable (Rios et al., 2018; Wulfmeier et al., 2017). To stabilize the training process, we use the concept of AC-GAN (Odena et al., 2017). AC-GAN uses an auxiliary task layer on top of the discriminator and appears to stabilize the adversarial learning procedure and improve the performance (Odena et al., 2017).

Following this concept, for the second task of  $D(\cdot)$ , we introduce an MRC-module  $D_q$  similar to  $M_Q$  on top of  $D_e$  as an auxiliary task layer.  $D_q$  enforces that the discriminator does not lose task-specific information while learning to encode domain-variant features. Later in Sections 4.4 and 4.5.2, we demonstrate the effectiveness of the auxiliary layer by performing an ablation study and a stability analysis. The input of  $D_q$  is the output of  $D_e$  for  $f_i^s$  and the output is the probability distributions for the start and end token positions of the answer span, similar to  $(p_i^{start}, p_i^{end})$ . Thus, the loss function  $\mathcal{L}_{D,aux}$

for  $D_q$  is the same as  $\mathcal{L}_Q$ . The final loss function  $\mathcal{L}_D$  for  $D(\cdot)$  is shown in Equation (5).

$$\mathcal{L}_D = \mathcal{L}_{D,triplet} + \mathcal{L}_{D,aux}. \quad (5)$$

### 3.2.4 Cost function

To eliminate domain shift by learning domain-invariant features, we integrate  $M_F(\cdot)$ ,  $M_Q(\cdot)$ , and  $D(\cdot)$  into adversarial learning framework, where we update  $M_F$  and  $M_Q$  to maximize  $\mathcal{L}_D$  and minimize  $\mathcal{L}_Q$  while updating  $D$  to minimize  $\mathcal{L}_D$ . Thus, the cost function  $\mathcal{L}_{total}$  of the BioADAPT-MRC framework consists of  $\mathcal{L}_Q$  and  $\mathcal{L}_D$  as shown in Equation (6) and is optimized end-to-end:

$$\mathcal{L}_{total} = \mathcal{L}_Q - \lambda \mathcal{L}_D. \quad (6)$$

Here,  $\lambda$  is a regularization parameter to balance  $\mathcal{L}_Q$  and  $\mathcal{L}_D$ . Unlike the original adversarial learning framework proposed in GAN, where the adversaries are updated alternately (Goodfellow et al., 2014), we perform joint optimization for all three components of our model using the gradient-reversal layer (Ganin and Lempitsky, 2015), as suggested by Chen et al. (2018a).

## 4 Results and discussion

We perform an extensive study to evaluate the proposed framework and compare with the SOTA biomedical-MRC methods on a collection of publicly available and widely used benchmark biomedical-MRC datasets.

### 4.1 Dataset

To demonstrate the effectiveness of our framework, we evaluate BioADAPT-MRC and compare it with the SOTA methods on three biomedical-MRC datasets from the BioASQ annual challenge (Tsatsaronis et al., 2015). The BioASQ competition has been organized since 2013 and consists of two large-scale biomedical NLP tasks, one of which is question answering (task B). Among four types of questions in task B, the factoid questions resemble the extractive biomedical-MRC task. As such, we utilize only the factoid MRC data from the BioASQ challenges held in 2019 (BioASQ-7b), 2020 (BioASQ-8b) and 2021 (BioASQ-9b) as the target-domain datasets to verify our model. These datasets were created from the search engine for biomedical literature, PubMed, with the help of domain-experts. Note that, for training, our framework requires only unlabeled contexts in the target domain. As such, we only consider the contexts in the BioASQ-7b, 8b and 9b training sets and disregard the question-answer pairs. The details on the availability of the training data and the pre-processing steps are provided in the Supplementary Section S1. At test time, we use the golden enriched test sets—BioASQ-7b, 8b and 9b—from the BioASQ challenges.

Similar to the previous studies (Jeong et al., 2020; Yoon et al., 2020), as the source-domain dataset, we use SQuAD-1.1 (Rajpurkar et al., 2016), which was developed from Wikipedia articles by crowd-workers. Table 1 shows the basic statistical information of all datasets used in the experiments. As shown, the number of training data samples in the source domain is noticeably higher than that of the target domain. The details on experimental setup and training configurations are provided in the Supplementary Section S2.

### 4.2 Metrics

For evaluation, we use three metrics used in the MRC task in the official BioASQ challenge: strict accuracy (SAcc), lenient accuracy (LAcc) and mean reciprocal rank (MRR). The BioASQ challenge requires the participant systems to predict the five best-matched answer spans extracted from the

context(s) in a decreasing order based on confidence score. In the test set, for each question, the biomedical experts in the BioASQ team provided one golden answer extracted from the context. Both golden answers and predicted answer spans are used to calculate the SAcc, LAcc and MRR scores, as shown in Equation (7). SAcc shows the models' capability to find exact answer location, LAcc determines the models' understanding of predicted answers' range and MRR reflects the quality of the predicted answer spans (Tsatsaronis et al., 2015):

$$SAcc = \frac{c_1}{n_{test}}; LAcc = \frac{c_5}{n_{test}}; MRR = \frac{1}{n_{test}} \sum_{i=1}^{n_{test}} \frac{1}{r(i)}. \quad (7)$$

Here,  $c_1$  is the number of questions correctly answered by the predicted answer span with the highest confidence score,  $c_5$  is the number of questions answered correctly by any of the five predicted answer spans,  $n_{test}$  is the number of questions in the test set and  $r(i)$  is the rank of the golden answer among all five predicted answer spans for the  $i$ th question. If the golden answer does not belong to the five predicted answer spans, we consider  $\frac{1}{r(i)} = 0$ . We implement these metrics by leveraging the publicly available tools provided by the official BioASQ challenge at <https://github.com/BioASQ/Evaluation-Measures>.

### 4.3 Method comparison

We compare the test-time performance of BioADAPT-MRC on BioASQ-7b and 8b with six best-performing models selected based on related published articles: Google (Hosein et al., 2019), BioBERT (Yoon et al., 2020), UNCC (Telukuntla et al., 2019), Umass (Kommaraju et al., 2020), KU-DMIS-2020 (Jeong et al., 2020) and BioQAExternalFeatures (Xu et al., 2021). For BioASQ-9b, we pick the best-performing system Ir\_sys2 from the BioASQ-9b leaderboard (available at: <http://participants-area.bioasq.org/results/9b/phaseB/>).

We also consider a hypothetical system that we would get for BioASQ-9b if that system would achieve the highest SAcc, LAcc and MRR scores on the leaderboard in all five batches of this test set. Note that, in reality, no individual system in the competition achieved the highest scores for all three metrics in all the batches (see Supplementary Section S3 and Supplementary Table S11 for details).

In addition to these models, we also compare the performance of BioADAPT-MRC with AdaMRC (Wang et al., 2019), a SOTA domain adaptation method for the MRC task. We provide brief descriptions of these aforementioned models in Supplementary Section S3.

### 4.4 Experimental results

Table 2 shows the comparison of BioADAPT-MRC with the SOTA biomedical-MRC methods on BioASQ-7b, 8b and 9b. As shown, BioADAPT-MRC improves on both LAcc and MRR when tested on all three BioASQ test sets and achieves the best performance. We also notice that while our model achieves the highest SAcc score for BioASQ-9b, it achieves the second-best SAcc scores for BioASQ-7b and 8b. The higher SAcc and LAcc scores imply that our model is able to correctly extract complete answers from the given contexts more frequently than the previous methods. The higher MRR scores, on the other hand, reflect our model's ability to extract complete answers with higher probability than the previous methods. In contrast to the previous works, our method uses no label information (question-answer pairs) during the training process and has still been able to achieve good performance, implying the effectiveness of our proposed framework.

As explained in Section 3, in the framework, we propose a domain-similarity discriminator with an auxiliary task layer that aims at promoting the generation of domain-invariant features in the feature extractor and thus improving the performance of the model. To show the effectiveness of the discriminator and the auxiliary task layer, we perform an ablation study and report the experimental results in Table 3. For a fair comparison, we perform all experiments under the same hyper-parameter settings. The baseline model shown in Table 3 consists of only the feature extractor and the MRC-module and was trained on the labeled source-domain dataset, SQuAD. For the remaining two models, we use the labeled SQuAD and the unlabeled BioASQ training datasets simultaneously. The addition of the discriminator enables the feature extractor in the baseline model to use the unlabeled BioASQ training datasets for generating domain-invariant feature representations. This is achieved by using the dissimilarity measurements between the feature representations of the SQuAD and BioASQ training data. As shown, after adding

**Table 1.** Statistics of the datasets used in the experiments

Dataset name	Training set (raw)	Training set (pre-processed)	Target to source ratio in training set	Test set
SQuAD-1.1	87 599	87 599	—	—
BioASQ-7b	779	5537	~1:16	162
BioASQ-8b	941	10 147	~1:9	151
BioASQ-9b	1092	13 178	~1:7	163

**Table 2.** Performance of BioADAPT-MRC compared with the best scores on BioASQ-7b, BioASQ-8b and BioASQ-9b test sets

Model	BioASQ-7b			BioASQ-8b			BioASQ-9b		
	SAcc	LAcc	MRR	SAcc	LAcc	MRR	SAcc	LAcc	MRR
Google (Hosein et al., 2019)	0.4201	0.5822	0.4798	—	—	—	—	—	—
BioBERT (Yoon et al., 2020)	0.4367	0.6274	0.5115	—	—	—	—	—	—
UNCC (Telukuntla et al., 2019)	0.3554	0.4922	0.4063	—	—	—	—	—	—
Umass (Kommaraju et al., 2020)	—	—	—	0.3133	0.4798	0.3780	—	—	—
KU-DMIS-2020 (Jeong et al., 2020)	<b>0.4510</b>	0.6245	0.5163	0.3819	0.5719	0.4593	—	—	—
BioQAExternalFeatures (Xu et al., 2021)*	0.4444	<i>0.6419</i>	<i>0.5165</i>	<b>0.3937</b>	<i>0.6098</i>	<i>0.4688</i>	—	—	—
BioASQ-9b Challenge—Best system (Ir_sys2)	—	—	—	—	—	—	0.5031	0.6626	0.5667
BioASQ-9b Challenge—Hypothetical system	—	—	—	—	—	—	<b>0.5399</b>	<i>0.7300</i>	<i>0.6017</i>
AdaMRC (Wang et al., 2019)	0.4321	0.6235	0.5136	0.3510	0.5828	0.4455	0.5337	0.7117	0.6001
BioADAPT-MRC	<i>0.4506</i>	<b>0.6420</b>	<b>0.5289</b>	<i>0.3841</i>	<b>0.6225</b>	<b>0.4749</b>	<b>0.5399</b>	<b>0.7423</b>	<b>0.6187</b>

Note: The best and the second-best scores are respectively highlighted in bold and italic.

‘—’ indicates that the corresponding source did not report the scores. \*denotes previously best-performing method for BioASQ-7B and BioASQ-8B.

**Table 3.** Test scores for ablation experiments of BioADAPT-MRC

Model	BioASQ-7b			BioASQ-8b			BioASQ-9b		
	SAcc	LAcc	MRR	SAcc	LAcc	MRR	SAcc	LAcc	MRR
Baseline	0.4136	0.6296	0.5056	0.3642	0.5960	0.4602	0.5092	0.7362	0.6010
BioADAPT-MRC (no auxiliary layer)	<i>0.4259</i>	<i>0.6296</i>	<i>0.5146</i>	<i>0.3775</i>	<i>0.6026</i>	<i>0.4679</i>	<i>0.5276</i>	<b>0.7485</b>	<i>0.6142</i>
BioADAPT-MRC	<b>0.4506</b>	<b>0.6420</b>	<b>0.5289</b>	<b>0.3841</b>	<b>0.6225</b>	<b>0.4749</b>	<b>0.5399</b>	<i>0.7423</i>	<b>0.6187</b>

Note: The best and the second-best scores are respectively highlighted in bold and italic.

only the discriminator without the auxiliary task layer, the performance of the model improves from the baseline, suggesting the influence of the discriminator. We explain this influence on the feature extractor more elaborately later in Section 4.5.6. For the final experiment in the ablation study (Table 3), we use our whole model consisting of the domain-similarity discriminator with the auxiliary task layer and notice an even further performance improvement. The auxiliary task layer, in this study, constrains the changes in the task-relevant features in the domain-similarity discriminator during training. Thus, the improvement in model performance after incorporating the auxiliary task layer suggests that with the task layer, the domain-similarity discriminator can better promote the generation of domain-invariant features that are simultaneously discriminative from the viewpoint of the MRC task in the source domain. Moreover, as explained in Section 3.2.3, we further demonstrate the stabilizing capability of the auxiliary task layer in Section 4.5.2.

## 4.5 Analysis

In this section, we analyze different components of the proposed framework. We also study the domain adaptation capability and the strengths and weaknesses of the BioADAPT-MRC model.

### 4.5.1 Triplet versus distance-based loss

BioADAPT-MRC uses triplet loss to optimize the discriminator. As explained in Section 3.2.3, we also consider using distance-based loss in place of triplet loss.

Table 4 shows the results of the contrastive experiments of loss functions—triplet loss and distance-based loss. As shown in Table 4, the model with triplet loss outperforms the one with the distance-based loss with higher mean SAcc, LAcc, MRR and lower standard deviation.

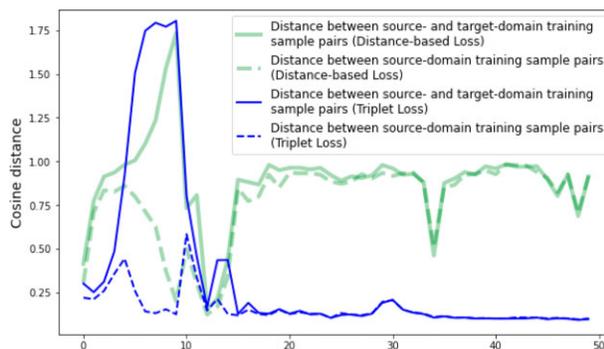
To further analyze this performance gap, we examine the trend of distance between source-domain training sample pairs and between source- and target-domain training sample pairs per epoch across 50 training epochs (Fig. 2).

Note that both experiments in Figure 2 are performed under the same seed. As shown in Figure 2, when we use distance-based loss, the cosine distance between either source-domain training sample pairs or source- and target-domain training sample pairs tends to be higher than when we use the triplet loss. It may happen because in the adversarial framework, while

**Table 4.** Average test scores with standard deviations (across three different seeds) for the contrastive experiments of discriminator loss functions—triplet loss and distance-based loss

Discriminator loss	BioASQ-9b		
	SAcc	LAcc	MRR
Distance-based	0.5256 ± 0.0202	0.7219 ± 0.0104	0.6038 ± 0.0105
Triplet	<b>0.5358 ± 0.0029</b>	<b>0.7321 ± 0.0077</b>	<b>0.6140 ± 0.0035</b>

Note: The best scores are highlighted in bold.

**Fig. 2.** Per-epoch cosine distance between source-domain training sample pairs and source- and target-domain training sample pairs across 50 epochs

distance-based loss focuses only on minimizing the distance between the source- and target-domain training samples without considering the distance between the source samples, triplet loss focuses on balancing both (Chen et al., 2017; Wang et al., 2017). As a result, triplet loss can minimize the domain shift to a greater extent than the distance-based loss and thus enable our framework to achieve higher performance.

**Table 5.** Average test scores with standard deviations across experiments with three random seeds (10, 42 and 2018) for initialization, to measure and compare the stability of BioADAPT-MRC

Model	BioASQ-7b		
	SAcc	LAcc	MRR
AdaMRC	<i>0.4300 ± 0.0029</i>	0.6152 ± 0.0116	<i>0.5083 ± 0.0076</i>
Baseline	0.4156 ± 0.0127	0.6173 ± 0.0101	0.5038 ± 0.0033
BioADAPT-MRC (no auxiliary layer)	0.4177 ± 0.0058	<i>0.6193 ± 0.0105</i>	0.5038 ± 0.0076
BioADAPT-MRC	<b>0.4465 ± 0.0029</b>	<b>0.6379 ± 0.0029</b>	<b>0.5237 ± 0.0037</b>
Model	BioASQ-8b		
	SAcc	LAcc	MRR
AdaMRC	0.3422 ± 0.0083	0.5960 ± 0.0143	0.4425 ± 0.0031
Baseline	0.3554 ± 0.0125	0.5960 ± 0.0162	0.4547 ± 0.0152
BioADAPT-MRC (no auxiliary layer)	<i>0.3664 ± 0.0113</i>	<i>0.5982 ± 0.0031</i>	<i>0.4618 ± 0.0080</i>
BioADAPT-MRC	<b>0.3797 ± 0.0031</b>	<b>0.6137 ± 0.0083</b>	<b>0.4750 ± 0.0024</b>
Model	BioASQ-9b		
	SAcc	LAcc	MRR
AdaMRC	0.5276 ± 0.0050	0.7239 ± 0.0100	0.6068 ± 0.0062
Baseline	0.5174 ± 0.0058	0.7198 ± 0.0126	0.6018 ± 0.0011
BioADAPT-MRC (no auxiliary layer)	<i>0.5337 ± 0.0050</i>	<i>0.7280 ± 0.0153</i>	<i>0.6127 ± 0.0012</i>
BioADAPT-MRC	<b>0.5358 ± 0.0029</b>	<b>0.7321 ± 0.0077</b>	<b>0.6140 ± 0.0035</b>

Note: The best and the second-best scores are respectively highlighted in bold and italic.

**Table 6.** Average test scores with standard deviations (across three different seeds) for experiments using synthetic questions and masked questions in the target-domain training dataset

Questions	BioASQ-9b		
	SAcc	LAcc	MRR
Synthetic	<i>0.5337 ± 0.0087</i>	<i>0.7198 ± 0.0126</i>	<i>0.6069 ± 0.0059</i>
Masked	<b>0.5358 ± 0.0029</b>	<b>0.7321 ± 0.0077</b>	<b>0.6140 ± 0.0035</b>

Note: The best and the second-best scores are respectively highlighted in bold and italic.

#### 4.5.2 Stability analysis

We examine the stability of BioADAPT-MRC and perform an error analysis of its performance by repeating the experiments for three different random seeds (10, 42 and 2018). Table 5 shows the test scores averaged across three seeds with standard deviations for BioASQ-7b, 8b and 9b. We also report the average scores with standard deviations for our baseline, the BioADAPT-MRC model with no auxiliary layer, and a SOTA method AdaMRC to compare model stability. As shown in Table 5, BioADAPT-MRC outperforms the other models with lower standard deviations, indicating higher stability of our framework. Moreover, the scores from the BioADAPT-MRC models with and without the auxiliary layer indicate that the auxiliary task layer helps increase both performance and overall model stability.

#### 4.5.3 Masked versus synthetic questions

Recall that, BioADAPT-MRC uses a special token [MASK] in place of the question tokens  $Q$  for the unlabeled target-domain training samples. The [MASK] tokens are used to inform the encoder model that the question tokens are missing and maintain consistency in the structure of the tokenized samples. Another approach to address the issue of missing questions in the target-domain training samples is to use synthetic questions (Wang et al., 2019). In Table 6, we present the results of the contrastive experiments of these two approaches—masked and synthetic questions. Inspired by the success of the AdaMRC question-generator in various target domains, such as news, Wikipedia and web search log, we use it to generate the synthetic questions.

Table 6 shows the average test scores for BioASQ-9b with standard deviations across three seeds. We find that although the synthetic questions can

noticeably improve performance over the baseline (see results for ‘Baseline’ in Table 5), BioADAPT-MRC with synthetic questions is unable to achieve better performance than with masked questions. It may happen because the biomedical domain differs from other domains, such as news or Wikipedia, in many linguistic dimensions, such as syntax, lexicon and semantics (Lee et al., 2020; Verspoor et al., 2009). As a result, while the question-generator can generate meaningful questions for domains, such as web, news and movie reviews (Wang et al., 2019), it mostly generates incoherent questions for the biomedical domain (as shown in Supplementary Fig. S4), which can eventually hurt the performance of the model. Moreover, synthetic-question generation also requires additional computational time—generating questions from around 10 000 contexts using a trained question-generator took ~3h with our computational resources (see configurations in Supplementary Section S2). Given the findings, we think that generating synthetic questions for the biomedical domain requires more attention and consider it as a future study.

#### 4.5.4 Semi-supervised setting

As an additional experiment, we evaluate the BioADAPT-MRC framework under a semi-supervised setting, where we combine labeled and unlabeled target-domain training data. We perform four experiments where the ratios of labeled target-domain training samples are 0%, 10%, 50% and 80% of the total target-domain training data. Note that we choose the labeled target-domain data by random sampling. Supplementary Table S5 shows the test scores on BioASQ-9b when trained with varying ratios of labeled target-domain data. As shown, with increased ratio of labeled samples in the target-domain training data, the performance scores also increase, which is expected. These results suggest that our proposed framework is also effective in a semi-supervised setting.

Note that although multiple labeled datasets in various sub-domains of biomedical-MRC (such as scientific literature and clinical notes) have been made available in the past few years (Pampari et al., 2018; Tsatsaronis et al., 2015), there is still a severe scarcity of labeled data in some other sub-domains that are linguistically different (e.g. consumer health biomedical-MRC) (Jin et al., 2023; Nguyen, 2019).

#### 4.5.5 Results on emrQA

We further validate our framework on another type of biomedical-MRC dataset, emrQA (Pampari et al., 2018), built using unstructured textual electronic health records (EHRs) with questions reflecting the inquiries made by clinicians about patients’ EHRs. The dataset contains five subsets, three of which

are extractive MRC datasets—heart disease risk, relations and medications. For our experiments, we use the heart disease risk dataset as the target-domain dataset. We randomly sample 10% of the dataset as the test set. To measure the performance, we use the widely used metrics for the extractive MRC task: Exact Match (EM) and *F1*-score (Baradaran et al., 2020). We compare the test scores with our baseline and the AdaMRC model. Supplementary Table S6 shows that BioADAPT-MRC improves the performance scores over baseline (9.67% in EM and 10.86% in *F1*) and AdaMRC (2.08% in EM and 2.16% in *F1*). This experiment validates that BioADAPT-MRC can be applied to different types of biomedical-MRC datasets.

We want to emphasize the fact that researchers have identified the unstructured clinical notes as inherently noisy and long with long-term textual dependencies (Cohen et al., 2013; Mahub et al., 2022; Pampari et al., 2018). We suspect that these phenomena may lead to an overall low EM and *F1* score (Joshi et al., 2017). Hence, we think that achieving higher scores in an MRC task on EHRs requires additional and rigorous data pre-processing and leave it as a future work.

#### 4.5.6 Domain adaptation

We show the influence of the domain-similarity discriminator by plotting (Supplementary Fig. S7) all samples from the BioASQ-9b test set and a set of random samples from the SQuAD training set. We pick random samples from the SQuAD training set to match the number of samples in the BioASQ-9b test set. As explained in Section 3.2.1, we use the feature representation of the [CLS] token as an accumulated representation of the whole input sequence. Each feature representation of the [CLS] token has a dimension of 768. To reduce these dimensions into two for visualization, we use multidimensional scaling (MDS) (Kruskal, 1964). We use MDS because it reduces the dimensions by preserving the dissimilarities between two data points in the original high-dimensional space. Since we use cosine distance in the discriminator to measure the dissimilarity between two domains, as the dissimilarity measure in the MDS, we use the pairwise cosine distance. The feature representations of the [CLS] token on the left plot and the right plot in Supplementary Figure S7 are generated by the feature extractors from the baseline model and the BioADAPT-MRC model, respectively. For a fair comparison, the selection of random SQuAD training samples is the same for the baseline and BioADAPT-MRC models. As shown, the features generated by the baseline model create two separate clusters for SQuAD and BioASQ-9b. The features generated by the BioADAPT-MRC model, on the other hand, form two overlapping clusters implying the reduced dissimilarities between the source and target domains. Interestingly, we notice that the data points from the BioASQ are closer to its cluster than those from the SQuAD. It may be because, unlike SQuAD, the data in the BioASQ originate from one single domain, and thus the feature representations are more similar to one another.

To further analyze the quality of the clusters before and after introducing the domain-similarity discriminator to the framework and thus to quantify the effect of domain adaptation, we perform DBSCAN clustering (Ester et al., 1996). We perform clustering on the MDS components of the features for the [CLS] tokens for the samples in the BioASQ test sets and the random samples from the SQuAD training set. Considering the bias of random sampling, for each BioASQ test set, we select five sets of random samples from the SQuAD training set and report the mean accuracy and silhouette scores with standard deviation in Supplementary Table S9 and Figure 3. We use the DBSCAN clustering because it views clusters as high-density regions where the distance between the samples is measured by a distance metric, providing flexibility in shapes and numbers of clusters. We describe the selected hyperparameters for the DBSCAN algorithm implementation in Supplementary Section S8.

Supplementary Table S9 and Figure 3 show that DBSCAN can identify two clusters with high accuracy when the features of the samples are extracted from the baseline model. The accuracy goes down when the features of the same samples are extracted from the BioADAPT-MRC model as they form a single cluster. Moreover, we analyze the silhouette scores to understand the separation distance between clusters. The range of silhouette score is  $[-1, 1]$ . A score of one indicates that the clusters are highly dense and clearly distinguishable from each other whereas  $-1$  refers to incorrect clustering. A score of zero or near zero indicates indistinguishable or overlapping clusters. As shown in Supplementary Table S9 and Figure 3, in this case, the high silhouette scores (closer to one) for the baseline model reflect that the feature representations of the samples from the same domain are very similar to its own cluster compared to the other one. On contrary, the low silhouette scores (closer to zero) for the BioADAPT-MRC model indicate that the feature

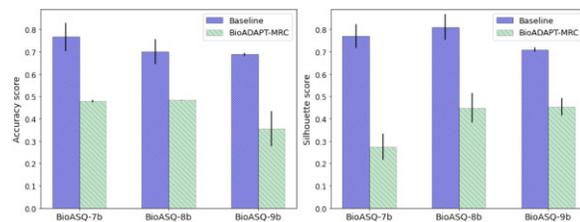


Fig. 3. Mean accuracy scores (left) and mean silhouette scores (right) with standard deviations for DBSCAN clustering on BioASQ test sets and SQuAD

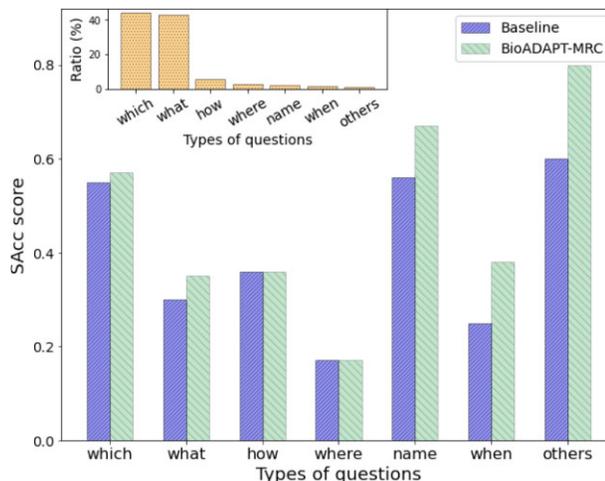


Fig. 4. Error analysis of BioADAPT-MRC, in comparison with the baseline model, depending on the types of questions in the BioASQ test sets

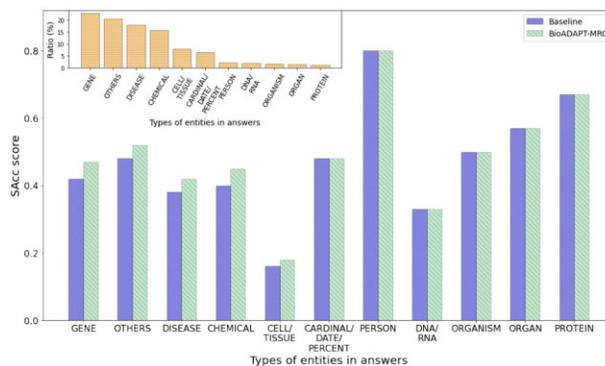


Fig. 5. Error analysis of BioADAPT-MRC, in comparison with the baseline model, depending on the types of answers in the BioASQ test sets

representations of the samples from both domains are very similar to one another. These results show the effectiveness of the domain-similarity discriminator in the BioADAPT-MRC framework. Considering the variability of the predicted answers in an MRC task, we present a motivating example to demonstrate how the word importance may impact the answer predictions and thus the performance of the biomedical-MRC task (see Supplementary Section S10). The example shows the effectiveness of BioADAPT-MRC over the baseline model for the given sample.

#### 4.5.7 Error analysis

We analyze the strengths and weaknesses of our approach by performing 2-fold error analysis. We focus on two aspects of the MRC dataset—types of questions and answers. Through this error analysis, we aim to answer the following questions: (i) what types of questions can be answered after domain adaptation that could not be answered before? (ii) What types of answers can be identified by the proposed approach after domain adaptation? (iii) What

<p><b>Example 1:</b> Question: What is the most common monogenic cause of common variable immunodeficiency (CVID) in Europeans? Golden Answer: <b>Heterozygous loss-of-function variants in NFKB1</b> Baseline Model: <b>NFKB1</b> (probability: 0.4902) BioADAPT-MRC Model: <b>Heterozygous loss-of-function variants in NFKB1</b> (probability: 0.5527)</p>	<p><b>Example 3:</b> Question: Which disease is treated with Anti-Siglec-8 Antibody? Golden Answer: <b>Eosinophilic Gastritis and Duodenitis</b> Baseline Model: <b>Eosinophilic Gastritis</b> (probability: 0.3602) BioADAPT-MRC Model: <b>Eosinophilic Gastritis and Duodenitis</b> (probability: 0.8733)</p>	<p><b>Example 5:</b> Question: What antibiotic is currently used as the standard of care for Clostridium Difficile infection as of 2018? Golden Answer: <b>Fidaxomicin</b> Baseline Model: <b>Fidaxomicin</b> (probability: 0.8742) BioADAPT-MRC Model: <b>Fidaxomicin</b> (probability: 0.9727)</p>	<p><b>Example 7:</b> Question: What is anophthalmia? Golden Answer: <b>absence of one or both eyes</b> Baseline Model: <b>autosomal dominant syndrome caused by mutations in the SOX2 gene and clinically characterized by severe eye malformations</b> (probability: 0.3139) BioADAPT-MRC Model: <b>severe eye malformations</b> (probability: 0.4540)</p>
<p><b>Example 2:</b> Question: What does "28" stand for in the Disease Activity Score DAS28? Golden Answer: <b>28 joints</b> Baseline Model: <b>subjective components</b> (probability: 0.4456) BioADAPT-MRC Model: <b>28 joints</b> (probability: 0.4951)</p>	<p><b>Example 4:</b> Question: What impacts stability of genomic imprinting in mouse pluripotent stem cells? Golden Answer: <b>A Susceptibility Locus on Chromosome 13</b> Baseline Model: <b>A Susceptibility Locus</b> (probability: 0.4693) BioADAPT-MRC Model: <b>A Susceptibility Locus on Chromosome 13</b> (probability: 0.5629)</p>	<p><b>Example 6:</b> Question: Approximately how many genes are contained in the X chromosome's non-pseudoautosomal region (non-PAR)? Golden Answer: <b>783</b> Baseline Model: <b>783</b> (probability: 0.8582) BioADAPT-MRC Model: <b>783</b> (probability: 0.9869)</p>	<p><b>Example 8:</b> Question: What is the indication for KYMRIAH? Golden Answer: <b>children and young adults with refractory or relapse B cell precursor acute lymphoblastic leukemia</b> Baseline Model: <b>refractory or relapse (R/R) B cell precursor acute lymphoblastic leukemia</b> (probability: 0.5414) BioADAPT-MRC Model: <b>treatment of pediatric patients and young adults with refractory or relapse (R/R) B cell precursor acute lymphoblastic leukemia</b> (probability: 0.4813)</p>

Fig. 6. Example question–answer pairs from the test sets demonstrating the strengths and weaknesses of the BioADAPT-MRC model over the baseline model. The green and red colors show correctly and incorrectly predicted answers, respectively (A color version of this figure appears in the online version of this article.)

types of items does the model struggle with, even with the domain adaptation component?

We use the SAcc score to perform the error analysis since the ultimate goal for any MRC system is to predict correct answer spans with the highest probability, reflected in the SAcc score. We categorize the test sets based on the following types of questions: *which*, *what*, *how*, *where*, *when* and *name* (an example question for *name* – ‘Name a CFL2 mutation which is associated with nemaline myopathy?’). We find that the most prevalent question types in the test sets are *what* (43%) and *which* (44%) (embedded in Fig. 4). We also find that after adding the domain adaptation module, the SAcc score increases by 52%, 20%, 17% and 4% for question types—*when*, *name*, *what* and *which*, respectively (Fig. 4). It indicates that the domain adaptation module can help increase the model’s capability to answer these types of questions with higher probability. For question types *how* and *where*, we do not notice further improvement, even with the domain adaptation module.

To analyze the types of answers that can be identified after introducing the domain adaptation module, we categorize the test sets based on the named entities in the answers. We use NER algorithms from the widely used spaCy library (Honnibal and Montani, 2017).

Figure 5 shows that the answers in the test sets mainly consist of entities, such as *GENE* (23%), *DISEASE* (18%), *CHEMICAL* (16%), *CELL/TISSUE* (8%) and *CARDINAL/DATE/PERCENT* (7%). As shown in Figure 5, after adding the domain adaptation module, the model has been able to identify the *CHEMICAL*, *CELL/TISSUE* entities with the highest SAcc scores (improvement by 13% from the baseline). We also notice an improvement in the SAcc scores for *GENE* and *DISEASE* entities. However, we notice no improvement in the SAcc scores for *CARDINAL/DATE/PERCENT*, *PERSON*, *DNA/RNA*, *ORGANISM*, *ORGAN* and *PROTEIN* entities even after adding the domain adaptation module. Given the results, we would like to emphasize that in both of the aforementioned analyses, for some categories, we do not have a large enough sample set to draw a definite conclusion about the effectiveness of the domain adaptation module and hence requires further future investigation.

To provide further evidence about what the model learned, we present eight example question–answer pairs demonstrating the strengths and weaknesses of the BioADAPT-MRC model over the baseline model (Fig. 6).

We randomly select two examples from each of the following four categories: (i) mispredicted by baseline, correctly predicted by BioADAPT-MRC, (ii) incompletely predicted by baseline, correctly predicted by BioADAPT-MRC, (iii) correctly predicted by both baseline and BioADAPT-MRC and (iv) mispredicted by both baseline and BioADAPT-MRC. In Examples 1–8, the answer spans respectively contain *GENE*, *CARDINAL*, *DISEASE*, *CELL*, *CHEMICAL*, *CARDINAL*, *ORGAN* and *DISEASE* entities. As shown in Examples 1–6, BioADAPT-MRC model can identify the answer span better than the baseline model with higher probability score. However, the probability scores (i.e. the prediction capability) can be further improved. Moreover, Examples 7 and 8 provide additional motivation for future investigation of the reason behind the misprediction of the model.

The results from the overall error analysis indicate that while the BioADAPT-MRC model does well under various scenarios, there is still significant room for potential improvement.

## 5 Conclusion

Biomedical-MRC is a crucial and emerging task in the biomedical domain pertaining to NLP. Biomedical-MRC aims at perceiving complex contexts

from the biomedical domain and helping medical professionals to extract information from them. Most MRC methods rely on a high volume of human-annotated data for near or similar to human-level performance. However, acquiring a labeled MRC dataset in the biomedical domain is expensive in terms of domain expertise, time and effort, creating the need for transfer learning from a source domain to a target domain. Due to variance between two domains, directly transferring an MRC model to the target domain often negatively affects its performance. We propose a framework for biomedical-MRC, BioADAPT-MRC, addressing the issue of domain shift by using a domain adaptation technique in an adversarial learning setting. We use a labeled MRC dataset from a general-purpose domain (source domain) along with unlabeled contexts from the biomedical domain (target domain) as our training data. We introduce a domain-similarity discriminator, aiming to reduce the domain shift between the general-purpose domain and biomedical domain to help boost the performance of the biomedical-MRC model. We validate our proposed framework on three widely used benchmark datasets from the biomedical question answering and semantic indexing challenge, BioASQ. We comprehensively demonstrate that without any label information in the target domain during training, the BioADAPT-MRC framework can achieve SOTA performance on these datasets. We perform an extensive quantitative study on the domain adaptation capability using dimensionality reduction and clustering techniques and show that our framework can learn domain-invariant feature representations. Additionally, we extend our framework to a semi-supervised setting and demonstrate that our framework can be efficiently applied even with varying ratios of labeled data. We perform a 2-fold error analysis to investigate the shortcomings of our framework and provide motivation for further future investigation and improvement.

We conclude that BioADAPT-MRC may be beneficial in healthcare systems as a tool to efficiently retrieve information from complex narratives and thus save valuable time and effort of the healthcare professionals.

The following are some future research directions that can originate from this work: (i) developing a synthetic question–answer generator specializing in the biomedical domain. (ii) Focusing on rigorous data pre-processing for the MRC task on unstructured clinical notes. (iii) Performing further investigation on the cases where BioADAPT-MRC struggles to improve over the baseline model. (iv) Applying BioADAPT-MRC to other NLP applications in the biomedical domain that suffer from labeled-data-scarcity issues. Such applications are biomedical NER, clinical negation detection, etc. (v) Analyzing the robustness of the domain-invariant feature representations learned by the BioADAPT-MRC model against meticulously crafted adversarial attack scenarios that may leverage syntactic and lexical knowledge-base from the dataset.

## Acknowledgements

We thank the anonymous reviewers for their insightful comments and suggestions.

## Author contributions

Conception, design, implementations, analysis, data pre-processing, interpretation, result analysis and original draft: M.M. Review and editing: M.M., S.S., E.B. and G.D.P. Supervision: E.B. and G.D.P.

## Funding

This work was supported by Department of Veterans Affairs, VHA Office of Mental Health and Suicide Prevention. This work has been authored by UT-Battelle, LLC under Contract No. [DE-AC05-00OR22725] with the US Department of Energy.

This research used resources of the Knowledge Discovery Infrastructure at the Oak Ridge National Laboratory, which is supported by the Office of Science of the US Department of Energy under Contract No. [DE-AC05-00OR22725]; and the Department of Veterans Affairs Office of Information Technology Inter-Agency Agreement with the Department of Energy under IAA No. [VA118-16-M-1062].

*Conflict of interest:* The authors declare that there is no conflict of interest.

## Disclaimer

The views and opinions expressed in this manuscript are those of the authors and do not necessarily represent those of the Department of Veterans Affairs, Department of Energy, or the United States Government. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

## Data availability

The data underlying this article are available in the BioASQ challenge website, at <http://participants-area.bioasq.org/datasets/>, and in the emrQA GitHub repository, at <https://github.com/panushri25/emrQA#download-dataset>.

## References

- Agnikula Kshatriya, B.S. et al. (2021) Identification of asthma control factor in clinical notes using a hybrid deep learning model. *BMC Med. Inform. Decis. Mak.*, 21, 1–10.
- Alsentzer, E. et al. (2019) Publicly available clinical BERT embeddings. In: Proceedings of the 2nd Clinical Natural Language Processing Workshop, Minneapolis, Minnesota, USA, Association for Computational Linguistics, pp. 72–78.
- Baradaran, R. et al. (2020) A survey on machine reading comprehension systems. *Natural Language Engineering*, First View, pp. 1–50.
- Bashath, S. et al. (2022) A data-centric review of deep transfer learning with applications to text data. *Inf. Sci.*, 585, 498–528.
- Bastian, H. et al. (2010) Seventy-five trials and eleven systematic reviews a day: how will we ever keep up? *PLoS Med.*, 7, e1000326.
- Bojanowski, P. et al. (2017) Enriching word vectors with subword information. *TACL*, 5, 135–146.
- Bromley, J. et al. (1993) Signature verification using a “siamese” time delay neural network. In: *Advances in Neural Information Processing Systems*. Vol. 6.
- Chen, W. et al. (2017) Beyond triplet loss: a deep quadruplet network for person re-identification. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA*, pp. 403–412.
- Chen, X. et al. (2018a) Adversarial deep averaging networks for cross-lingual sentiment classification. *TACL*, 6, 557–570.
- Chen, Y. et al. (2018b) DarkRank: accelerating deep metric learning via cross sample similarities transfer. In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 32.
- Clark, K. et al. (2020) ELECTRA: pre-training text encoders as discriminators rather than generators. In: *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- Cohen, R. et al. (2013) Redundancy in electronic health record corpora: analysis, impact on text mining performance and mitigation strategies. *BMC Bioinformatics*, 14, 10–15.
- Del Fiol, G. et al. (2014) Clinical questions raised by clinicians at the point of care: a systematic review. *JAMA Intern. Med.*, 174, 710–718.
- Devlin, J. et al. (2019) BERT: pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Minneapolis, Minnesota, USA, Association for Computational Linguistics, pp. 4171–4186.
- Ester, M. et al. (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. *KDD*, 96, 226–231.
- Fox, S. and Duggan, M. (2013) Health online 2013. *Health*, 2013, 1–55.
- Ganin, Y. and Lempitsky, V. (2015) Unsupervised domain adaptation by back-propagation. In: *International Conference on Machine Learning, Lille, France*. PMLR, pp. 1180–1189.
- Glorot, X. et al. (2011) Domain adaptation for large-scale sentiment classification: a deep learning approach. In: *Proceedings of the 28th International Conference on Machine Learning, Bellevue, WA, USA, 2011*.
- Golub, D. et al. (2017) Two-stage synthesis networks for transfer learning in machine comprehension. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 835–844. Association for Computational Linguistics, Copenhagen, Denmark.
- Goodfellow, I. et al. (2014) Generative adversarial nets. In: *Advances in Neural Information Processing Systems, Montréal, Canada*. Vol. 27.
- Gu, Y. et al. (2021) Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans. Comput. Healthcare (HEALTH)*, 3, 1–23.
- Guan, D. et al. (2021) Scale variance minimization for unsupervised domain adaptation in image segmentation. *Pattern Recognit.*, 112, 107764.
- Hider, P.N. et al. (2009) The information-seeking behavior of clinical staff in a large health care organization. *J. Med. Libr. Assoc.*, 97, 47–50.
- Honnibal, M. and Montani, I. (2017) SpaCy 2: natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. *Appear*, 7, 411–420.
- Hosein, S. et al. (2019) Measuring domain portability and error propagation in biomedical QA. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Würzburg, Germany*, pp. 686–694. Springer.
- Jeong, M. et al. (2020) Transferability of natural language inference to biomedical question answering. In: *Working Notes of CLEF 2020 Conference and Labs of the Evaluation Forum, Thessaloniki, Greece*.
- Jin, Q. et al. (2023) Biomedical question answering: a survey of approaches and challenges. *ACM Comput. Surv.*, 55, 1–36.
- Joshi, M. et al. (2017) TriviaQA: a large scale distantly supervised challenge dataset for reading comprehension. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611. Association for Computational Linguistics, Vancouver, Canada.
- Kaya, M. and Bilge, H.Ş. (2019) Deep metric learning: a survey. *Symmetry*, 11, 1066.
- Kim, W. et al. (2018) Attention-based ensemble for deep metric learning. In: *Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germany*, pp. 736–751.
- Kommaraju, V. et al. (2020) Unsupervised pre-training for biomedical question answering. In: *CLEF (Working Notes), Thessaloniki, Greece*.
- Kruskal, J.B. (1964) Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29, 1–27.
- Laiz, P. et al. (2019) Using the triplet loss for domain adaptation in WCE. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops, Seoul, Korea (South)*.
- Lee, J. et al. (2020) BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36, 1234–1240.
- Mahub, M. et al. (2022) Unstructured clinical notes within the 24 hours since admission predict short, mid & long-term mortality in adult ICU patients. *PLoS One*, 17, e0262182.
- Mikolov, T. et al. (2013) Efficient estimation of word representations in vector space. In: Bengio, Y. and LeCun, Y. (eds) 1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2–4, 2013, *Workshop Track Proceedings*.
- Naseem, U. et al. (2021) BioALBERT: a simple and effective pre-trained language model for biomedical named entity recognition. In: *2021 International Joint Conference on Neural Networks (IJCNN), Online*, pp. 1–7. IEEE.
- Nguyen, V. (2019) Question answering in the biomedical domain. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop, Florence, Italy*, pp. 54–63.
- Nishida, K. et al. (2020) Unsupervised domain adaptation of language models for reading comprehension. In: *LREC, Marseille, France*.
- Odena, A. et al. (2017) Conditional image synthesis with auxiliary classifier GANs. In: *International Conference on Machine Learning, Sydney, Australia*, PMLR, pp. 2642–2651.
- Pampari, A. et al. (2018) emrQA: a large corpus for question answering on electronic medical records. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for

- Computational Linguistics, Brussels, Belgium, pp. 2357–2368, Brussels, Belgium.
- Pan,S.J. and Yang,Q. (2010) A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.*, **22**, 1345–1359.
- Pennington,J. et al. (2014) GloVe: global vectors for word representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, pp. 1532–1543.
- Pergola,G. et al. (2021) Boosting low-resource biomedical QA via entity-aware masking strategies. In: Merlo,P. et al. (eds) *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, April 19–23, 2021, Online*, Association for Computational Linguistics, pp. 1977–1985.
- Raj Kanakarajan,K. et al. (2021) BioELECTRA: pretrained biomedical text encoder using discriminators. In: *Proceedings of the 20th Workshop on Biomedical Language Processing*, pp. 143–154.
- Rajpurkar,P. et al. (2016) SQuAD: 100,000+ questions for machine comprehension of text. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, Austin, Texas*, Association for Computational Linguistics, pp. 2383–2392.
- Rios,A. et al. (2018) Generalizing biomedical relation classification with neural adversarial domain adaptation. *Bioinformatics*, **34**, 2973–2981.
- Russell-Rose,T. and Chamberlain,J. (2017) Expert search strategies: the information retrieval practices of healthcare information professionals. *JMIR Med. Inform.*, **5**, e7680.
- Sun,S. et al. (2017) An unsupervised deep domain adaptation approach for robust speech recognition. *Neurocomputing*, **257**, 79–87.
- Telukuntla,S.K. et al. (2019) UNCC biomedical semantic question answering systems. BioASQ: Task-7B, Phase-B. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Würzburg, Germany*, Springer, pp. 695–710.
- Thompson,B. et al. (2019) Overcoming catastrophic forgetting during domain adaptation of neural machine translation. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Minneapolis, Minnesota, pp. 2062–2068.
- Tsatsaronis,G. et al. (2015) An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics*, **16**, 1–28.
- Tzeng,E. et al. (2017) Adversarial discriminative domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, Hawaii*, pp. 7167–7176.
- Verspoor,K. et al. (2009) The textual characteristics of traditional and open access scientific journals are similar. *BMC Bioinformatics*, **10**, 183.
- Vu,T.-T. et al. (2020) Effective unsupervised domain adaptation with adversarially trained language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online, Association for Computational Linguistics, pp. 6163–6173.
- Wang,H. et al. (2019) Adversarial domain adaptation for machine reading comprehension. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Association for Computational Linguistics, Hong Kong, China, pp. 2510–2520.
- Wang,J. et al. (2017) Deep metric learning with angular loss. In: *Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy*, pp. 2593–2601.
- Weinberger,K.Q. et al. (2005) Distance metric learning for large margin nearest neighbor classification. In: *Advances in Neural Information Processing Systems*. Vol. **18**, Vancouver, British Columbia, Canada.
- Wen,G. et al. (2018) Improving face recognition with domain adaptation. *Neurocomputing*, **287**, 45–51.
- Wulfmeier,M. et al. (2017) Addressing appearance change in outdoor robotics with adversarial domain adaptation. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vancouver, British Columbia, Canada, IEEE, pp. 1551–1558.
- Xu,G. et al. (2021) External features enriched model for biomedical question answering. *BMC Bioinformatics*, **22**, 1–19.
- Yoon,W. et al. (2020) Pre-trained language model for biomedical question answering. In: , and , *Machine Learning and Knowledge Discovery in Databases - International Workshops of ECML PKDD 2019, Proceedings*, Würzburg, Germany, Springer, pp. 727–740.
- Yue,X. et al. (2021) CliniQ4QA: generating diverse questions for domain adaptation of clinical question answering. In: *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, Online, pp. 580–587.