

Neural Domain Adaptation for Biomedical Question Answering

Georg Wiese^{1,2}, Dirk Weissenborn² and Mariana Neves¹

¹ Hasso Plattner Institute, August Bebel Strasse 88, Potsdam 14482 Germany

² Language Technology Lab, DFKI, Alt-Moabit 91c, Berlin, Germany

georg.wiese@student.hpi.de,

dewe01@dfki.de, mariana.neves@hpi.de

Abstract

Factoid question answering (QA) has recently benefited from the development of deep learning (DL) systems. Neural network models outperform traditional approaches in domains where large datasets exist, such as SQuAD ($\approx 100,000$ questions) for Wikipedia articles. However, these systems have not yet been applied to QA in more specific domains, such as biomedicine, because datasets are generally too small to train a DL system from scratch. For example, the BioASQ dataset for biomedical QA comprises less than 900 factoid (single answer) and list (multiple answers) QA instances. In this work, we adapt a neural QA system trained on a large open-domain dataset (SQuAD, *source*) to a biomedical dataset (BioASQ, *target*) by employing various transfer learning techniques. Our network architecture is based on a state-of-the-art QA system, extended with biomedical word embeddings and a novel mechanism to answer list questions. In contrast to existing biomedical QA systems, our system does not rely on domain-specific ontologies, parsers or entity taggers, which are expensive to create. Despite this fact, our systems achieve state-of-the-art results on factoid questions and competitive results on list questions.

1 Introduction

Question answering (QA) is the task of retrieving *answers* to a *question* given one or more *contexts*. It has been explored both in the open-domain setting (Voorhees et al., 1999) as well as domain-specific settings, such as BioASQ for the

biomedical domain (Tsatsaronis et al., 2015). The BioASQ challenge provides ≈ 900 *factoid* and *list* questions, i.e., questions with one and several answers, respectively. This work focuses on answering these questions, for example: *Which drugs are included in the FEC-75 regimen?* \rightarrow *fluorouracil, epirubicin, and cyclophosphamide*.

We further restrict our focus to *extractive* QA, i.e., QA instances where the correct answers can be represented as spans in the contexts. Contexts are relevant documents which are provided by an information retrieval (IR) system.

Traditionally, a QA pipeline consists of named-entity recognition, question classification, and answer processing steps (Jurafsky, 2000). These methods have been applied to biomedical datasets, with moderate success (Zi et al., 2016). The creation of large-scale, open-domain datasets such as SQuAD (Rajpurkar et al., 2016) have recently enabled the development of neural QA systems, e.g., Wang and Jiang (2016), Xiong et al. (2016), Seo et al. (2016), Weissenborn et al. (2017), leading to impressive performance gains over more traditional systems.

However, creating large-scale QA datasets for more specific domains, such as the biomedical, would be very expensive because of the need for domain experts, and therefore not desirable. The recent success of deep learning based methods on open-domain QA datasets raises the question whether the capabilities of trained models are transferable to another domain via *domain adaptation* techniques. Although domain adaptation has been studied for traditional QA systems (Blitzer et al., 2007) and deep learning systems (Chen et al., 2012; Ganin et al., 2016; Bousmalis et al., 2016; Riemer et al., 2017; Kirkpatrick et al., 2017), it has to our knowledge not yet been applied for end-to-end neural QA systems.

To bridge this gap we employ various do-

main adaptation techniques to transfer knowledge from a trained, state-of-the-art neural QA system (FastQA, Weissenborn et al. (2017)) to the biomedical domain using the much smaller BioASQ dataset. In order to answer *list questions* in addition to factoid questions, we extend FastQA with a novel answering mechanism. We evaluate various transfer learning techniques comprehensively. For factoid questions, we show that mere fine-tuning reaches state-of-the-art results, which can further be improved by a forgetting cost regularization (Riemer et al., 2017). On list questions, the results are competitive to existing systems. Our manual analysis of a subset of the factoid questions suggests that the results are even better than the automatic evaluation states, revealing that many of the "incorrect" answers are in fact synonyms to the gold-standard answer.

2 Related Work

Traditional Question Answering Traditional factoid and list question answering pipelines can be subdivided into named-entity recognition, question classification, and answer processing components (Jurafsky, 2000). Such systems have also been applied to biomedical QA such as the OAQA system by Zi et al. (2016). Besides a number of domain-independent features, they incorporate a rich amount of biomedical resources, including a domain-specific parser, entity tagger and thesaurus to retrieve concepts and synonyms. A logistic regression classifier is used both for question classification and candidate answer scoring. For candidate answer generation, OAQA employs different strategies for general factoid/list questions, choice questions and quantity questions.

Neural Question Answering Neural QA systems differ from traditional approaches in that the algorithm is not subdivided into discrete steps. Instead, a single model is trained end-to-end to compute an answer directly for a given question and context. The typical architecture of such systems (Wang and Jiang, 2016; Xiong et al., 2016; Seo et al., 2016) can be summarized as follows:

1. **Embedding Layer:** Question and context tokens are mapped to a high-dimensional vector space, for example via GloVe embeddings (Pennington et al., 2014) and (optionally) character embeddings (Seo et al., 2016).
2. **Encoding Layer:** The token vectors are processed independently for question and context, usually by a recurrent neural network (RNN).
3. **Interaction Layer:** This layer allows for interaction between question and context representations. Examples are Match-LSTM (Wang and Jiang, 2016) and Coattention (Xiong et al., 2016).
4. **Answer Layer:** This layer assigns start and end scores to all of the context tokens, which can be done either statically (Wang and Jiang, 2016; Seo et al., 2016) or by a dynamic decoding process (Xiong et al., 2016).

FastQA *FastQA* fits into this schema, but reduces the complexity of the architecture by removing the interaction layer, while maintaining state-of-the-art performance (Weissenborn et al., 2017). Instead of one or several interaction layers of RNNs, FastQA computes two simple word-in-question features for each token, which are appended to the embedding vectors before the encoding layer. We chose to base our work on this architecture because of its state-of-the-art performance, faster training time and reduced number of parameters.

Unsupervised Domain Adaptation Unsupervised domain adaptation describes the task of learning a predictor in a *target domain* while labeled training data only exists in a different *source domain*. In the context of deep learning, a common method is to first train an autoencoder on a large unlabeled corpus from both domains and then use the learned input representations as input features to a network trained on the actual task using the labeled source domain dataset (Glorot et al., 2011; Chen et al., 2012). Another approach is to learn the hidden representations directly on the target task. For example, domain-adversarial training optimizes the network such that it computes hidden representations that both help predictions on the source domain dataset and are indistinguishable from hidden representations of the unlabeled target domain dataset (Ganin et al., 2016). These techniques cannot be straightforwardly applied to the question answering task, because they require a large corpus of biomedical question-context pairs (albeit no answers are required).

Supervised Domain Adaptation In contrast to the unsupervised case, supervised domain adaptation assumes access to a small amount of labeled training data in the target domain. The simplest approach to supervised domain adaptation for neural models is to pre-train the network on data from the source domain and then fine-tune its parameters on data from the target domain. The main drawback of this approach is *catastrophic forgetting*, which describes the phenomenon that neural networks tend to "forget" knowledge, i.e., its performance in the source domain drops significantly when they are trained on the new dataset. Even though we do not directly aim for good performance in the source domain, measures against catastrophic forgetting can serve as a useful regularizer to prevent over-fitting.

Progressive neural networks combat this issue by keeping the original parameters fixed and adding new units that can access previously learned features (Rusu et al., 2016). Because this method adds a significant amount of new parameters which have to be trained from scratch, it is not well-suited if the target domain dataset is small. Riemer et al. (2017) use fine-tuning, but add an additional *forgetting cost* term that punishes deviations from predictions with the original parameters. Another approach is to add an L2 loss which punishes deviation from the original parameters. Kirkpatrick et al. (2017) apply this loss selectively on parameters which are important in the source domain.

3 Model

Our network architecture is based on FastQA (Weissenborn et al., 2017), a state-of-the-art neural QA system. Because the network architecture itself is exchangeable, we treat it as a black box, with subtle changes at the input and output layer as well as to the decoding and training procedure. These changes are described in the following. See Figure 3 for an overview of the system.

3.1 Input Layer

In a first step, words are embedded into a high-dimensional vector space. We use three sources of embeddings, which are concatenated to form a single embedding vector:

- **GloVe embeddings:** 300-dimensional GloVe vectors (Pennington et al., 2014). These are

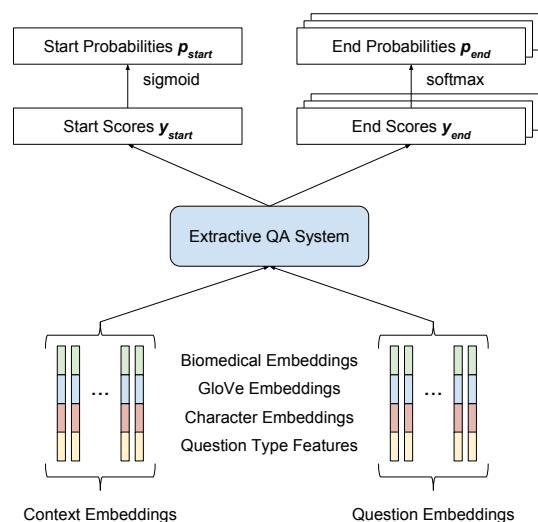


Figure 1: Network architecture of our system for biomedical question answering. At its core, it uses an extractive neural QA system as a black box (we use FastQA (Weissenborn et al., 2017)). The embedding layer is modified in order to include biomedical word embeddings and question type features. The output layer is adjusted to add the ability to answer list questions in addition to factoid questions.

open-domain word vectors trained on 840 billion tokens from web documents. The vectors are not updated during training.

- **Character embeddings:** As used in FastQA (Weissenborn et al., 2017) and proposed originally by Seo et al. (2016), we employ a 1-dimensional convolutional neural network which computes word embeddings from the characters of the word.
- **Biomedical Word2Vec embeddings:** 200-dimensional vectors trained using Word2Vec (Mikolov et al., 2013) on about 10 million PubMed abstracts (Pavlopoulos et al., 2014). These vectors are specific to the biomedical domain and we expect them to help on biomedical QA.

As an optional step, we add entity tag features to the token embeddings via concatenation. Entity tags are provided by a dictionary-based entity tagger based on the UMLS Metathesaurus. The entity tag feature vector is a 127-dimensional bit vector that for each of the UMLS semantic types states whether the current token is part of an entity of that type. This step is only applied if explicitly

noted.

Finally, a one-hot encoding of the question type (*factoid* or *list*) is appended to all the input vectors. With these embedding vectors as input, we invoke FastQA to produce start and end scores for each of the n context tokens. We denote start scores by y_{start}^i and end scores conditioned on a predicted start at position i by $y_{end}^{i,j}$, with start index $i \in [1, n]$ and end index $j \in [i, n]$.

3.2 Output Layer

In our adapted output layer, we convert the start and end scores to span probabilities. The computation of these probabilities is independent of the question type. The interpretation, however, depends on the question type: While for factoid questions, the list of answer spans is interpreted as a ranked list of answer candidates, for list questions, answers above a certain probability threshold are interpreted as the set of answers to the question.

Given the start scores $y_{start}^1, \dots, y_{start}^n$ and end scores $y_{end}^{i,1}, \dots, y_{end}^{i,n}$, we compute the start and end probabilities as follows:

$$p_{start}^i = \sigma(y_{start}^i) \quad (1)$$

$$p_{end}^{i,\cdot} = \text{softmax}(y_{end}^{i,\cdot}) \quad (2)$$

where $\sigma(x)$ is the sigmoid function. As a consequence, multiple tokens can be chosen as likely start tokens, but the network is expected to select a single end token for a given start token, hence the softmax function. Finally, the probability that a given span (i, j) answers the question is $p_{span}^{i,j} = p_{start}^i \cdot p_{end}^{i,j}$. This extension generalizes the FastQA output layer such that multiple answer spans with different start positions can have a high probability, allowing us to retrieve multiple answers for list questions.

3.3 Decoding

Given a trained model, start probabilities can be obtained by running a forward pass and computing the start probability as in Equation 1. For the top 20 starts, we compute the end probabilities as given by Eq. 2. From the start and end probabilities, we extract the top 20 answer spans ranked by $p_{span}^{i,j}$. As a simple post-processing step, we remove duplicate strings and retain only those with the highest probability.

For factoid questions, we output the 5 most likely answer spans as our ranked list of answers. For list questions, we learn a *probability cutoff threshold* t that defines the set of list answers $A = \{(i, j) | p_{span}^{i,j} \geq t\}$. We choose t to be the threshold that optimizes the list F1 score on the respective development set.

3.4 Domain Adaptation

Fine-tuning Our training procedure consists of two phases: In the *pre-training phase*, we train the model on SQuAD, using a token F1 score as the training objective as by Weissenborn et al. (2017). We will refer to the resulting parameters as the *base model*. In the *fine-tuning phase*, we initialize the model parameters with the base model and then continue our optimization on the BioASQ dataset with a smaller learning rate.

Forgetting Cost Regularization To avoid *catastrophic forgetting* during fine-tuning as a means to regularize our model, we optionally add an additional forgetting cost term L_{fc} , as proposed by Riemer et al. (2017). It is defined as the cross-entropy loss between the current predictions and the base model’s predictions.

L2 Weight Regularization We also add an L2 loss term L_{l2} which penalizes deviations from the base model’s parameters. Note that a more advanced approach would be to apply this loss selectively on weights which are particularly important in the source domain (Kirkpatrick et al., 2017). The final loss is computed as $L_{final} = L_{original} + C_{fc} \cdot L_{fc} + C_{l2} \cdot L_{l2}$ where C_{fc} and C_{l2} are hyperparameters which are set to 0 unless otherwise noted.

4 Experimental Setup

4.1 Datasets

SQuAD *SQuAD* (Rajpurkar et al., 2016) is a dataset of $\approx 100,000$ questions with relevant contexts and answers that sparked research interest into the development of neural QA systems recently. The contexts are excerpts of Wikipedia articles for which crowd-source workers generated questions-answer pairs. Because of the large amount of training examples in SQuAD, it lends itself perfectly as our source dataset.

BioASQ The BioASQ challenge provides a biomedical QA dataset (Tsatsaronis et al., 2015)

consisting of questions, relevant contexts (called *snippets*) from PubMed abstracts and possible answers to the question. It was carefully created with the help of biomedical experts.

In this work, we focus on Task B, Phase B of the BioASQ challenge, in which systems must answer questions from gold-standard snippets. These questions can be either yes/no questions, summary questions, factoid questions, or list questions. Because we employ an extractive QA system, we restrict this study to answering factoid and list questions by extracting answer spans from the provided contexts.

The 2017 BioASQ training dataset contains 1,799 questions, of which 413 are factoid and 486 are list questions. The questions have ≈ 20 snippets on average, each of which are on average ≈ 34 tokens long. We found that around 65% of the factoid questions and around 92% of the list questions have at least one extractable answer. For questions with extractable answers, answer spans are computed via a simple substring search in the provided snippets. All other questions are ignored during training and treated as answered incorrectly during evaluation.

4.2 Training

We minimize the cross-entropy loss for the gold standard answer spans. However, for multiple answer spans that refer to the same answer (e.g. synonyms), we only minimize the loss for the span of the lowest loss. We use the ADAM (Kingma and Ba, 2014) for optimization on SQuAD with a learning rate starting at 10^{-3} which is halved whenever performance drops between checkpoints. During the fine-tuning phase, we continue optimization on the BioASQ dataset with a smaller learning rate starting at 10^{-4} . During both phases, the model is regularized by variational dropout of rate 0.5 (Gal and Ghahramani, 2015).

4.3 Evaluation

The official evaluation measures from BioASQ are mean reciprocal rank (MRR) for factoid questions and F1 score for list questions¹. For factoid questions, the list of ranked answers can be at most five entries long. The F1 score is measured on the gold standard list elements. For both measures,

¹The details can be found at http://participants-area.bioasq.org/Tasks/b/eval_meas/

case-insensitive string matches are used to check the correctness of a given answer. A list of synonyms is provided for all gold-standard answers. If the system's response matches one of them, the answer counts as correct.

For evaluation, we use two different fine-tuning datasets, depending on the experiment: *BioASQ3B*, which contains all questions of the first three BioASQ challenges, and *BioASQ4B* which additionally contains the test questions of the fourth challenge. *BioASQ4B* is used as the training dataset for the fifth BioASQ challenge whereas *BioASQ3B* was used for training during the fourth challenge.

Because the datasets are small, we perform 5-fold cross-validation and report the average performance across the five folds. We use the larger *BioASQ4B* dataset except when evaluating the ensemble and when comparing to participating systems of previous BioASQ challenges.

All models were implemented using TensorFlow (Abadi et al., 2016) with a hidden size of 100. Because the context in BioASQ usually comprises multiple snippets, they are processed independently in parallel for each question. Answers from all snippets belonging to a question are merged and ranked according to their individual probabilities.

5 Results

5.1 Domain Adaptation

In this section, we evaluate various domain adaptation techniques. The results of the experiments are summarized in Table 1.

Baseline As a baseline without transfer learning, Experiment 1 trains the model on BioASQ only. Because the BioASQ dataset by itself is very small, a dropout rate of 0.7 was used, because it worked best in preliminary experiments. We observe a rather low performance, which is expected when applying deep learning to such a small dataset.

Fine-tuning Experiments 2 and 3 evaluate the pure fine-tuning approach: Our *base model* is a system trained on SQuAD only and tested on BioASQ (Experiment 2). For Experiment 3, we fine-tuned the base model on the *BioASQ4B* training set. We observe that performance increases significantly, especially on list questions. This increase is expected, because the network is trained

Experiment	Factoid MRR	List F1
(1) Training on BioASQ only	17.9%	19.1%
(2) Training on SQuAD only	20.0%	8.1%
(3) Fine-tuning on BioASQ	24.6%	23.6%
(4) Fine-tuning on BioASQ w/o biomedical embeddings	21.3%	22.4%
(5) Fine-tuning on BioASQ w/ entity features	23.3%	23.8%
(6) Fine-tuning on BioASQ + SQuAD	23.9%	23.8%
(7) Fine-tuning on BioASQ w/ forgetting cost ($C_{fc} = 100.0$)	26.2%	21.1%
(8) Fine-tuning on BioASQ w/ L2 loss on original parameters ($C_{l2} = 0.3$)	22.6%	20.4%

Table 1: Comparison of various transfer learning techniques. In Experiment 1, the model was trained on BioASQ only. In Experiment 2, the model was trained on SQuAD and tested on BioASQ. We refer to it as the *base model*. In Experiment 3, the base model parameters were fine-tuned on the BioASQ training set. Experiments 4-5 evaluate the utility of domain dependent word vectors and features. Experiments 6-8 address the problem of catastrophic forgetting. All experiments have been conducted with the *BioASQ4B* dataset and 5-fold cross-validation.

on biomedical- and list questions, which are not part of the SQuAD dataset, for the first time. Overall, the performance of the fine-tuned model on both question types is much higher than the baseline system without transfer learning.

Features In order to evaluate the impact of using biomedical word embeddings, we repeat Experiment 3 without them (Experiment 4). We see a factoid and list performance drop of 3.3 and 1.2 percentage points, respectively, showing that biomedical word embeddings help increase performance.

In Experiment 5, we append entity features to the word vector, as described in Section 3.1. Even though these features provide the network with domain-specific knowledge, we found that it actually harms performance on factoid questions. Because most of the entity features are only active during fine-tuning with the small dataset, we conjecture that the performance decrease is due to over-fitting.

Catastrophic Forgetting We continue our study with techniques to combat catastrophic forgetting as a means to regularize training during fine-tuning. In Experiment 6 of Table 1 we fine-tune the base model on a half-half mixture of BioASQ and SQuAD questions (BioASQ questions have been upsampled accordingly). This form of joint training yielded no significant performance gains. Experiment 7 regularizes the model via an additional forgetting cost term, as

proposed by [Riemer et al. \(2017\)](#) and explained in Section 3.4. We generally found that this technique only increases performance for factoid questions where the performance boost was largest for $C_{fc} = 100.0$. The fact that the forgetting loss decreases performance on list questions is not surprising, as predictions are pushed more towards the predictions of the base model, which has very poor performance on list questions.

Experiment 8 adds an L2 loss which penalizes deviations from the base model’s parameters. We found that performance decreases as we increase the value of C_{l2} which shows that this technique does not help at all. For the sake of completeness we report results for $C_{l2} = 0.3$, the lowest value that yielded a significant drop in performance.

5.2 Ensemble

Model ensembles are a common method to tweak the performance of a machine learning system. Ensembles combine multiple model predictions, for example by averaging, in order to improve generalization and prevent over-fitting. We evaluate the utility of an ensemble by training five models on the *BioASQ3B* dataset using 5-fold cross-validation. Each of the models is evaluated on the 4B test data, i.e., data which is not included in *BioASQ3B*.

During application, we run an ensemble by averaging the start and end scores of individual models before they are passed to the sigmoid / softmax functions as defined in Eq. 1 and 2. In Table 2 we summarize the average performance of

Experiment	Factoid MRR	List F1
Average	23.4%	24.0%
Best	24.3%	27.7%
Ensemble	27.3%	28.6%

Table 2: Performance of a model ensemble. Five models have been trained on the *BioASQ3B* dataset and tested on the 4B test questions. We report the average and best single model performances, as well as the ensemble performance.

the five models, the best performance across the five models, and the performance of the ensemble. We observe performance gains of 3 percentage points on factoid questions and a less than 1 percentage point on list questions, relative to the best single model. This demonstrates a small performance gain that is consistent with the literature.

5.3 Comparison to competing BioASQ systems

Because the final results of the fifth BioASQ challenge are not available at the time of writing, we compare our system to the best systems in last year’s challenge². For comparison, we use the best single model and the model ensemble trained on *BioASQ3B* (see Section 5.2). We then evaluate the model on the 5 batches of last year’s challenge using the official BioASQ evaluation tool. Each batch contains 100 questions of which only some are factoid and list questions. Note that the results underestimate our system’s performance, because our competing system’s responses have been manually evaluated by humans while our system’s responses are evaluated automatically using string matching against a potentially incomplete list of synonyms. In fact, our qualitative analysis in Section 5.4 shows that many answers are counted as incorrect, but are synonyms of the gold-standard answer. The results are summarized in Table 3 and compared to the best systems in the challenge in each of the batches and question type categories.

With our system winning four out of five batches on factoid questions, we consider it state-of-the-art in biomedical factoid question answering, especially when considering that our results might be higher on manual evaluation. The results on list questions are slightly worse, but still very

²Last year’s results are available at <http://participants-area.bioasq.org/results/4b/phaseB/>

competitive. This is surprising, given that the network never saw a list question prior to the fine-tuning phase. Due to small test set sizes, the sampling error in each batch is large, causing the single model to outperform the model ensemble on some batches.

5.4 Qualitative Analysis

In order to get a better insight into the quality of the predictions, we manually validated the predictions for the factoid questions of batch 5 of the fourth BioASQ challenge as given by the best single model (see Table 3). There are in total 33 factoid questions, of which 23 have as the gold standard answer a span in one of the contexts. According to the official BioASQ evaluation, only 4 questions are predicted correctly (i.e., the gold standard answer is ranked highest). However, we identified 10 rank-1 answers which are not counted as correct but are synonyms to the gold standard answer. Examples include “CMT4D disease” instead of “Charcot-Marie-Tooth (CMT) 4D disease”, “tafazzin” instead of “Tafazzin (TAZ) gene”, and “ β -glucocerebrosidase” instead of “Beta glucocerebrosidase”. In total, we labeled 14 questions as correct and 24 questions as having their correct answer in the top 5 predictions.

In the following, we give examples of mistakes made by the system. Questions are presented in italics. In the context, we underline predicted answers and present correct answers in boldface.

We identified eight questions for which the semantic type of the top answer differs from the question answer type. Some of these cases are completely wrong predictions. However, this category also includes subtle mistakes like the following:

In which yeast chromosome does the rDNA cluster reside?

The rDNA cluster in Saccharomyces cerevisiae is located 450 kb from the left end and 610 kb from the right end of **chromosome XII**...

Here, it predicted a yeast species the rDNA cluster is located in, but ignored that the question is asking for a chromosome.

Another type of mistakes is that the top answer is somewhat correct, but is missing essential information. We labeled four predictions with this category, like the following example:

Batch	Factoid MRR			List F1		
	Best Participant	Single	Ensemble	Best Participant	Single	Ensemble
1	12.2% (fa1)	25.2%	29.2%	16.8% (fa1)	29.1%	27.9%
2	22.6% (LabZhu-FDU)	16.4%	24.2%	15.5% (LabZhu-FDU)	25.8%	20.8%
3	24.4% (oaqa-3b-3)	24.7%	20.6%	48.3% (oaqa-3b-3)	31.8%	33.3%
4	32.5% (oaqa-3b-4)	34.0%	40.3%	31.2% (oaqa-3b-4)	29.0%	24.1%
5	28.5% (oaqa-3b-5)	23.7%	23.2%	29.0% (oaqa-3b-5)	23.5%	26.1%
Avg.	24.0%	24.8%	27.5%	28.1%	27.8%	26.5%

Table 3: Comparison to systems on last year’s (fourth) BioASQ challenge for factoid and list questions. For each batch and question type, we list the performance of the best competing system, our single model and ensemble. Note that our qualitative analysis (Section 5.4) suggests that our factoid performance on batch 5 would be about twice as high if all synonyms were contained in the gold standard answers.

How early during pregnancy does non-invasive cffDNA testing allow sex determination of the fetus?

Gold Standard Answer: "6th to 10th week of gestation" or "first trimester of pregnancy"

Given Top Answer: "6th-10th"

In summary, to our judgment, 14 of 33 questions (42.4%) are answered correctly, and 24 of 33 questions (72.7%) are answered correctly in one of the top 5 answers. These are surprisingly high numbers considering low MRR score of 23.7% of the automatic evaluation (Table 3).

6 Discussion and future work

The most significant result of this work is that state-of-the-art results in biomedical question answering can be achieved even in the absence of domain-specific feature engineering. Most competing systems require structured domain-specific resources, such as biomedical ontologies, parsers, and entity taggers. While these resources are available in the biomedical domain, they are not available in most domains.

Our system, on the other hand, requires a large open-domain QA dataset, biomedical word embeddings (which are trained in an unsupervised fashion), and a small biomedical QA dataset. This suggests that our methodology is easily transferable to other domains as well.

Furthermore, we explored several supervised domain adaptation techniques. In particular, we demonstrated the usefulness of forgetting cost for factoid questions. The decreased performance on list questions is not surprising, because the model’s performance on those questions is very

poor prior to fine-tuning which is due to the lack of list questions in SQuAD. We believe that large scale open-domain corpora for list questions would enhance performance further.

Unsupervised domain adaptation could be an interesting direction for future work, because the biomedical domain offers large amounts of textual data, some of which might even contain questions and their corresponding answers. We believe that leveraging these resources holds potential to further improve biomedical QA.

7 Conclusion

In this paper, we described a deep learning approach to address the task of biomedical question answering by using domain adaptation techniques. Our experiments reveal that mere fine-tuning in combination with biomedical word embeddings yield state-of-the-art performance on biomedical QA, despite the small amount of in-domain training data and the lack of domain-dependent feature engineering. Techniques to overcome catastrophic forgetting, such as a forgetting cost, can further boost performance for factoid questions. Overall, we show that employing domain adaptation on neural QA systems trained on large-scale, open-domain datasets can yield good performance in domains where large datasets are not available.

Acknowledgments

This research was supported by the German Federal Ministry of Education and Research (BMBF) through Software Campus project GeNIE (01IS12050).

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.
- John Blitzer, Mark Dredze, Fernando Pereira, et al. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *ACL*. volume 7, pages 440–447.
- Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. 2016. Domain separation networks. In *Advances in Neural Information Processing Systems*. pages 343–351.
- Minmin Chen, Zhixiang Xu, Kilian Weinberger, and Fei Sha. 2012. Marginalized denoising autoencoders for domain adaptation. *arXiv preprint arXiv:1206.4683*.
- Yarin Gal and Zoubin Ghahramani. 2015. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. *arXiv preprint arXiv:1506.02142 2*.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *Journal of Machine Learning Research* 17(59):1–35.
- Xavier Glorot, Antoine Bordes, and Yoshua Bengio. 2011. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *Proceedings of the 28th international conference on machine learning (ICML-11)*. pages 513–520.
- Dan Jurafsky. 2000. *Speech & language processing*. Pearson Education India.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* page 201611835.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. pages 3111–3119.
- Ioannis Pavlopoulos, Aris Kosmopoulos, and Ion Androutsopoulos. 2014. [Continuous space word vectors obtained by applying word2vec to abstracts of biomedical articles](#) <http://bioasq.lip6.fr/info/BioASQword2vec/>.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. pages 1532–1543. <http://www.aclweb.org/anthology/D14-1162>.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Matthew Riemer, Elham Khabiri, and Richard Goodwin. 2017. Representation stability as a regularizer for improved text analytics transfer learning <https://openreview.net/pdf?id=HyenWc5gx>.
- Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. 2016. Progressive neural networks. *arXiv preprint arXiv:1606.04671*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- George Tsatsaronis, Georgios Balikas, Prodromos Malakasiotis, Ioannis Partalas, Matthias Zschunke, Michael R Alvers, Dirk Weissenborn, Anastasia Krithara, Sergios Petridis, Dimitris Polychronopoulos, et al. 2015. An overview of the bioasq large-scale biomedical semantic indexing and question answering competition. *BMC bioinformatics* 16(1):1.
- Ellen M Voorhees et al. 1999. The trec-8 question answering track report. In *Trec*. volume 99, pages 77–82.
- Shuohang Wang and Jing Jiang. 2016. Machine comprehension using match-1stm and answer pointer. *arXiv preprint arXiv:1608.07905*.
- Dirk Weissenborn, Georg Wiese, and Laura Seiffe. 2017. Making neural qa as simple as possible but not simpler. *arXiv preprint arXiv:1703.04816*.
- Caiming Xiong, Victor Zhong, and Richard Socher. 2016. Dynamic coattention networks for question answering. *arXiv preprint arXiv:1611.01604*.
- Yang Zi, Zhou Yue, and Eric Nyberg. 2016. Learning to answer biomedical questions: Oaqa at bioasq 4b. *ACL 2016* page 23.