

Query Focused Multi-document Summarisation of Biomedical Texts*

Macquarie University and the Australian National
University at BioASQ8b

Diego Mollá^[0000-0003-4973-0963]¹, Christopher Jones^[0000-0002-3491-739X]¹, and
Vincent Nguyen^[0000-0003-1787-8090]²

¹ Macquarie University, Australia

Diego.Molla-Aliod@mq.edu.au

Christopher.Jones4@hdr.mq.edu.au

² Australian National University, Australia

Vincent.Nguyen@anu.edu.au

Abstract. This paper presents the participation of Macquarie University and the Australian National University for Task B Phase B of the 2020 BioASQ Challenge (BioASQ8b). Our overall framework implements Query focused multi-document extractive summarisation by applying either a classification or a regression layer to the candidate sentence embeddings and to the comparison between the question and sentence embeddings. We experiment with variants using BERT and BioBERT, Siamese architectures, and reinforcement learning. We observe the best results when BERT is used to obtain the word embeddings, followed by an LSTM layer to obtain sentence embeddings. Variants using Siamese architectures or BioBERT did not improve the results.

1 Introduction

Query focused multi-document summarisation aims to generate the answer to a question by combining information from multiple documents [1]. This task, therefore, is related to both question answering and text summarisation. There is substantial research in both question answering and text summarisation. In the case of text summarisation, most research focuses on single-document summarisation, and there is also substantial research on multi-document summarisation. However, there is relatively little research on query focused multi-document text summarisation. There are multiple applications where query focused multi-document text summarisation can be useful. A clear example of a useful application is in the domain of biomedicine and clinical medicine, where a doctor or a

Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0). CLEF 2020, 22-25 September 2020, Thessaloniki, Greece.

* Code associated with this paper is available at <https://github.com/dmollaalioid/bioasq8b-public>

patient wants to obtain a concise summary of the most relevant evidence related to a particular diagnosis or intervention.

The BioASQ Challenge³ organises shared tasks centered on biomedical texts. The focus of this paper is on the 2020 participation of Macquarie University and the Australian National University in Task B Phase B (BioASQ8b), where the aim is to find the “ideal answer” to a question, given a collection of relevant PubMed abstracts.⁴ We approach this task as an instance of query focused multi-document extractive summarisation by scoring each candidate sentence and selecting the top-scoring ones to produce the final summary. Macquarie University and the Australian National University submitted independent runs to BioASQ8b, but we both used the starting code of Macquarie University’s BioASQ7b participation [10]. Novel contributions of this paper, compared with previous participation at BioASQ, include:

1. The incorporation of BERT and BioBERT in the general architecture of [10].
2. The use of Siamese architectures in two main setups: 1) sharing weights in the architecture of [10], and 2) using Sentence-BERT [11].
3. The use of Proximal Policy Optimisation (PPO) and BERT in a Reinforcement Learning approach.

This paper is structured as follows. Section 2 describes the general framework and baselines based on previous participation at BioASQ. Section 3 explains the incorporation of BERT and BioBERT. Section 4 introduces the incorporation of Siamese architectures. Section 5 describes the use of Reinforcement Learning. Section 6 presents and discusses the results of cross-validation evaluations using the BioASQ8b training data. Section 7 presents and discusses the runs submitted to BioASQ. Finally, Section 8 concludes this paper.

2 General Framework and Baselines

The overall architecture of our systems is based on that of Macquarie University’s participation to BioASQ7b [10]. This architecture is described in Fig. 1. The baseline systems re-use the following options from [10] to generate the embeddings of the words and sentences.

- The **embedding generator** to obtain word embeddings is a matrix of pre-trained embeddings generated by word2vec. We trained word2vec using PubMed documents provided by the organisers of BioASQ. The embeddings had a vector size of 100.
- The **embedding reductor** to obtain sentence embeddings is a pair of forward and backward LSTM chains. The weights of the embedding reductor for the candidate sentence and the question were not shared.
- The activation function of the **final layer** is a linear function for the regression setup and a sigmoid function for the classification setup.

³ <http://www.bioasq.org>

⁴ <http://www.ncbi.nlm.nih.gov/pubmed>

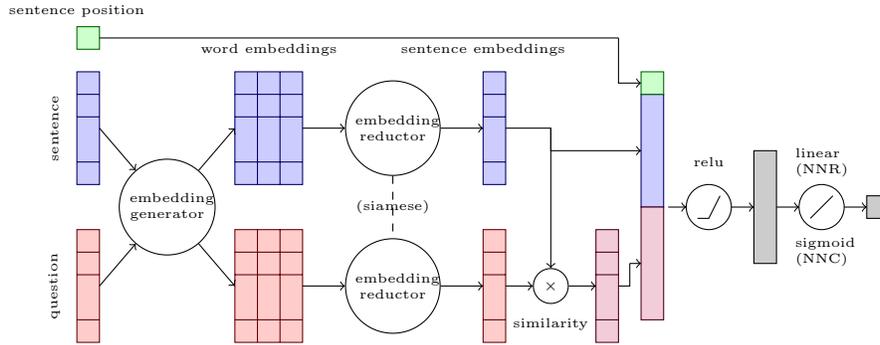


Fig. 1. Base architecture. An embedding generator calculates the embeddings of each of the words of the candidate sentence and the question. These are then passed to an embedding reducer that obtains the overall sentence embeddings of the candidate sentence and the question. A Siamese variant enforced shared weights of the embedding reducers. Then, the sentence position is concatenated with the sentence embedding and the similarity of sentence and question embeddings, implemented as a product. A final layer predicts the label of the sentence.

- The similarity between the embeddings of the candidate sentence and the question is the element-wise product.

As in [10], we tried a classification setup and a regression setup. In the regression setup (“NNR” in Fig. 1 Table 3), the training data is labelled with the ROUGE-SU4 F1 value of the candidate sentence and the objective function to optimise is the Mean Squared Error. In the classification setup (“NNC” in Fig. 1 Table 3), the 5 candidate sentences with the highest ROUGE-SU4 F1 are labelled as 1, and the rest are labelled as 0. The objective function to optimise in the classification setup is binary cross-entropy.

In both the classification and the regression setup, the summary is produced by scoring each candidate sentence and extracting the top n sentences to generate the summary, where n depended on the question type and was the same as reported in [10] (Table 1).

Table 1. Number of sentences selected, for each question type

	Summary Factoid Yesno List			
n	6	2	2	3

3 Experiments with BERT and BioBERT

BERT [2] has been used in a wide range of NLP tasks, including text classification [9], and extractive [8] and abstractive [7] summarisation. We have integrated BERT into our general architecture of Fig. 1 as described below.

In a first experiment (“BERT untrained” in Table 3), we replaced the embedding generator with BERT using the pre-trained model provided by Huggingface.⁵ The resulting word embeddings are now affected by context. The BERT weights were not updated during the training stage. Following the recommendation of the Huggingface library, the embedding reductor of each candidate sentence and the question are the average of the word embeddings.

In a subsequent experiment (“BERT trained” in Table 3), the embedding generator and reducers are as in BERT untrained, but we allowed the BERT weights to be fine-tuned during the training process.

We also tried a variant (“BERT LSTM” in Table 3) that uses BERT in the embedding generator as in BERT untrained, but the embedding reductor is a bidirectional LSTM chain as in the NNC baseline. This variant is therefore comparable with the NNC baseline, and the only difference being the use of BERT for the embedding generator.

In a final series of experiments (“BioBERT untrained” and “BioBERT LSTM” in Table 3), the embedding generator and reducers are as in BERT untrained and BERT LSTM, but the pre-trained model was as provided by the developers of BioBERT⁶ [6], who used biomedical documents to pre-train BERT.

In all of the experiments in this section, the final layer was a classification layer with a sigmoid activation and the objective function to optimise was binary cross-entropy.

4 Experiments with Siamese Networks

Siamese networks have been used in applications that include a comparison between documents, for example to determine semantic similarity [11]. The general idea is to use the same processing module for each of the two documents by sharing the weights of the encoding component that generates the embeddings of the documents. We have used this idea in two main kinds of experiments described below.

4.1 Siamese LSTM

A straightforward implementation of Siamese Networks (“Siamese LSTM” in Table 3) using the overall architecture of Fig. 1 shares the weights of the embedding reducers of the candidate sentence and the question. This ensures that the sentence embeddings are generated using exactly the same process.

⁵ <https://huggingface.co/>

⁶ <https://github.com/dmis-lab/biobert>

4.2 Sentence-BERT

The second implementation of Siamese Networks uses Sentence-BERT (SBERT) [11] to determine whether a candidate sentence is similar to the question. In particular, the system uses BERT [2] in a Siamese setup as described in Fig. 2.

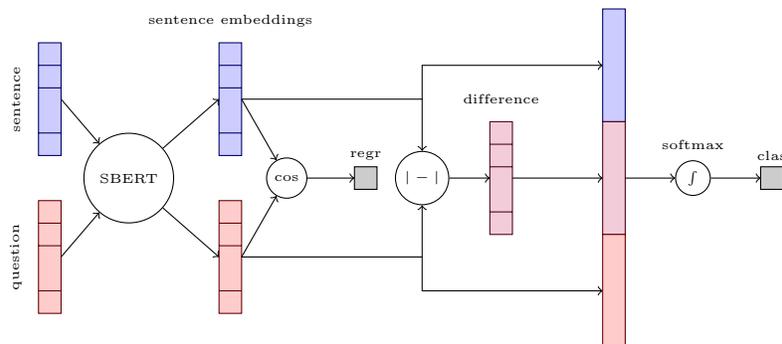


Fig. 2. Architecture for the experiments with SBERT. SBERT is used to generate the embeddings of the candidate sentence and the question. The embeddings are then processed in either a regression setup or a classification setup.

Depending on the settings, the system uses regression, classification or both (multi-task), for prediction.

- In the **regression** setup, the cosine similarity between candidate sentence and question embeddings is computed, and the objective function to optimise is the Mean Squared Error (MSE).
- In the **classification** setup, the candidate sentence and question embeddings are concatenated with the element-wise absolute difference between the sentence and question, and a softmax layer is added. The objective function to optimise is binary cross-entropy.
- In the **multi-task** setup, the classification and regression setups are jointly optimised during training. At prediction time, we use either the classification head or the regression head.

In both the regression and the classification setups, the training data was labelled with the classification labels mentioned in Section 2.

We also experimented with a BERT and a BioBERT variation⁷ of a Sentence Transformer [11] to ensure that the sentence embeddings reflect a biomedical word-space.

Table 2 shows all the combinations we tried.

⁷ <https://huggingface.co/gsarti/biobert-nli>

Table 2. Variants of SBERT used in our experiments

System Name	Model	Training		Prediction	
		Classification	Regression	Classification	Regression
SBERT R	BERT		Y		Y
SBERT C	BERT	Y		Y	
SBERT M R	BERT	Y	Y		Y
SBERT M C	BERT	Y	Y	Y	
SBioBERT R	BioBERT		Y		Y
SBioBERT C	BioBERT	Y		Y	
SBioBERT M R	BioBERT	Y	Y		Y
SBioBERT M C	BioBERT	Y	Y	Y	

5 Experiments with Reinforcement Learning

Reinforcement learning allows the training process to optimise the target evaluation metric (ROUGE-SU4 F1) directly. Whereas [10] used the REINFORCE algorithm [13], in our participation to BioASQ8b we used the Proximal Policy Optimisation (PPO) approach. We choose PPO for our summarisation task because past research shows that it penalises changes to the policy [12], and because we observe a more consistent learning curve using this approach compared to REINFORCE [5].

5.1 Approach

As in past submissions, our reinforcement learning system classifies sentences to be either (0) not included in the summary or (1) included in the summary based on a policy, using the ROUGE-SU4 F1 score directly as the reward. We use the stable baselines library⁸ [3] to implement the PPO reinforcement learning approach, and apply this to our BioASQ summarisation task environment.

We perform some hyperparameter tuning, while leaving the PPO code unmodified as much as possible. We use a horizon (n steps) of 1000 with 4 mini-batches each, and run for a total of 500,000 timesteps. When training our model, we choose any action from the probability distribution of the policy function, but when testing we choose the action out of 100 samples.

The neural network architecture for PPO is shown in Figure 3. As in [10], the inputs of the neural network consist of:

1. Candidate sentence
2. Question
3. Summary generated so far
4. Sentences after the candidate sentence
5. Entire document

⁸ <https://stable-baselines.readthedocs.io/>

6. Length of summary generated so far

For each of the first five inputs, we take the mean of the word2vec word embeddings (each of size 100) to generate the sentence embeddings. The sentence embeddings and the length of the summary are concatenated to form a single layer of size 501. We then feed the combined layer into a simple Multi-Layer Perceptron with two hidden layers each of size 200. The outputs of this neural network become the value function and stochastic policy function for our PPO approach, each of size 2 with linear activations, representing the predicted future rewards and action probabilities respectively for the two actions (classify 0 or 1).

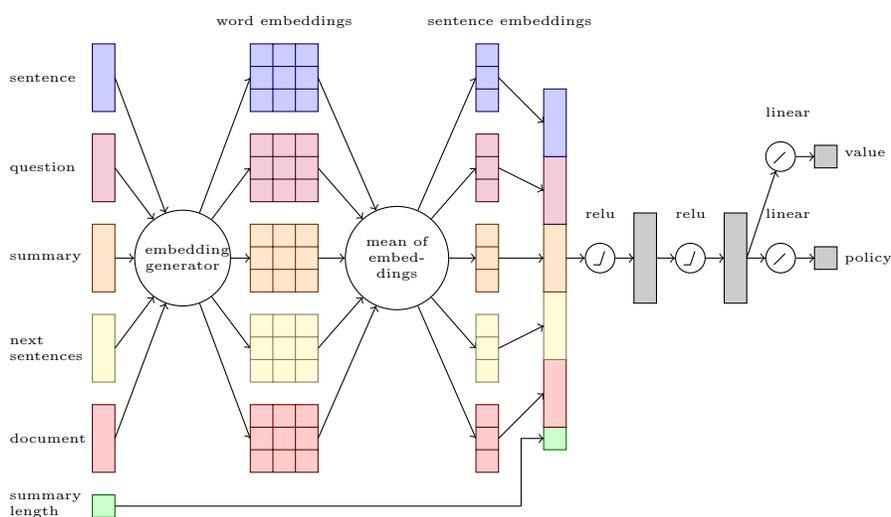


Fig. 3. Architecture for our PPO neural network. An embedding generator calculates the embeddings for each word, and the mean of all word embeddings is used as the sentence embedding. The 6 coloured inputs are combined into a single layer of size 501 and then fed into 2 hidden layers of size 200 each. The outputs are used as the policy function and value function for PPO.

5.2 With BERT

We also apply pre-trained BERT embeddings to our reinforcement learning approach for batches 4 and 5 of BioASQ8b. Our network architecture is the same as in Figure 3, but we change the embedding generator to generate BERT embeddings instead of word2vec embeddings. We use the PyTorch BERT embedding generator provided by Huggingface⁵ and the TensorFlow Multi-Layer Perceptron policy provided by stable baselines⁸. We observe a minor improvement in ROUGE-SU4 F1 score as shown in Section 7 (PPO BERT).

6 Results and Discussion

6.1 Cross-Validation Results

All experiments except those described in Sections 5 were evaluated using 10-fold cross-validation using the training data provided by the organisers of BioASQ8b. These results are shown in Table 3. As described in Section 6.2, cross-validation evaluation was not practical for the experiments with reinforcement learning.

Table 3. Results of experiments using ROUGE SU4 F-score under 10-fold cross-validation. The table shows the mean and standard deviation across the folds. “firstn” is a baseline that selects the first n sentences. NNR and NNC are described in Section 2. BERT and BioBERT are described in Section 3. Siamese LSTM is described in Section 4.1. The SBERT and SBioBERT runs are described in Section 4.2.

Method	ROUGE-SU4 F1 Mean \pm 1 stdev	
firstn	0.261 \pm 0.011	
NNR	0.264 \pm 0.008	
NNC	0.271 \pm 0.013	
BERT untrained	0.270 \pm 0.014	
BERT trained	0.261 \pm 0.012	
BERT LSTM	0.274 \pm 0.010	
BioBERT untrained	0.262 \pm 0.010	
BioBERT LSTM	0.264 \pm 0.012	
Siamese LSTM	0.263 \pm 0.010	

0.23 0.24 0.25 0.26 0.27 0.28

The baseline runs “firstn”, “NNR” and “NNC” confirm [10]’s observation that the classification setup produces better results than the regression setup. The untrained BERT system did not better the classification system, and the trained BERT system produced worse results.

We observed no changes in the evaluation results of the BERT trained system as we changed the number of epochs from 1 to 20 epochs. A detailed look at the changes in the loss during training revealed a very small improvement of the loss as we increase the number of epochs, but not enough to reflect a difference in the final ROUGE-SU4 F1 metric. This suggests that a more elaborate fine-tuning regime with gradual unfreezing as described by [4] might lead to improved results.

Given the poor results of the trained BERT system, we kept BERT frozen when we tested the variants using LSTM and using BioBERT.

BERT and BioBERT followed by an LSTM-based sentence reductor did improve results over the version with a mean of embeddings. Unfortunately, the experiments using the LSTM-based reductor were made after the deadlines for submission of results to BioASQ8b.

It was surprising to observe that the BioBERT variants performed worse than the BERT variants. This is not in line with the improvement in the performance of BioBERT for the “exact answers” section of BioASQ7b reported in literature [14].

The Siamese LSTM variant did not improve over the classification system. The reason for this might be, as mentioned in Section 7, that questions and candidate sentences are different in nature. We also conducted cross-validation experiments with the SBERT variants but they are not included in Table 3 as the results are very different from those of the runs submitted to BioASQ (Section 6) and we suspect that there might have been an error when running these evaluations.

Table 4 shows and explains the hyperparameters of all systems of Table 4.

Table 4. Hyperparameters for the experiments of Table 3. The choice of dropout and epochs is the result of grid search. A dropout of 0 means that no dropout was applied. The choice of batch size is determined by the GPU capabilities. In all cases, the dimensions of the word and sentence embeddings is 100. The size of the LSTM states and intermediate layers is also 100. The number of candidate sentences is limited to 50. The BERT version is bert-base-uncased provided by <https://huggingface.co/>. The BioBERT version is v1.1-pubmed available at <https://github.com/dmis-lab/biobert>

Method	Batch size	Dropout	Epochs	Sentence length
NNR	1024	0.3	10	300
NNC	1024	0.3	10	300
BERT untrained	32	0	50	250
BERT trained	8	0	1	250
BERT LSTM	1024	0.6	10	250
BioBERT untrained	1024	0	10	250
BioBERT LSTM	1024	0.6	10	250
Siamese LSTM	1024	0.2	10	300

6.2 Reinforcement Learning Results

The evaluation setup for reinforcement learning did not use cross-validation due to the long time it took to train the system (several days). We therefore used a partition of the training data for the train and test sets.

Table 5 shows the evaluation results of the PPO [12] reinforcement learning approaches discussed in Section 5, compared with the REINFORCE [13] approach, using the BioASQ7b and BioASQ8b training datasets.

Table 5. Reinforcement Learning Preliminary Results. The results using the BioASQ7b dataset show the mean and standard deviation of 3 runs. The results using the BioASQ8b dataset show the result of a single run. The PPO BERT system was only run with the BioASQ8b dataset and was only submitted to batches 4 and 5 of BioASQ8b. All results use the ROUGE-SU4 F1 metric.

System	BioASQ7b Mean \pm stdev	BioASQ8b
REINFORCE word2vec	0.253 \pm 0.001	0.265
PPO word2vec	0.251 \pm 0.001	0.269
PPO BERT	-	0.274

The **BioASQ7b** data set consists of 2,747 questions which we divide into a training set (2,289 questions) and a testing set (458 questions) partitioned 5:1 using a random seed to shuffle the data. The PPO word2vec and REINFORCE word2vec systems were run 3 times each for 500,000 timesteps on the BioASQ7b data set, and the maximum ROUGE-SU4 F1 score reached in each learning curve was averaged across the 3 runs.

The **BioASQ8b** data set consists of 3,243 questions which we divide into a training set (2,702 questions) and a testing set (541 questions) partitioned 5:1 using a random seed to shuffle the data. The PPO word2vec, PPO BERT, and REINFORCE word2vec systems were each run once only for 500,000 timesteps on the BioASQ8b data set, and the maximum ROUGE-SU4 F1 score for each run is shown in Table 5. The models which produced the maximum ROUGE-SU4 F1 score for each system were saved and re-used to generate our PPO submissions to the BioASQ competition in 2020, except for the REINFORCE system which was not included in this year’s submission.

Whereas our evaluation results did not show any clear difference between REINFORCE and PPO on word2vec embeddings, the version using PPO and BERT features showed an improvement.

The reinforcement learning experiments of Table 5, however, do not outperform the results of the experiments of Table 3.

7 Submissions to BioASQ8b

Table 6 shows the results of Macquarie University’s submissions to BioASQ8b. We only report ROUGE-SU4 F1 for comparison with our experiments. Also, as observed by [10], F1 scores have a higher correlation with human judgements than recall scores, and similar correlation with human judgements as precision scores.

Table 6. MQ runs submitted to BioASQ8b

Batch	Run Name	Description	ROUGE-SU4 F1
1	MQ1	First n	0.3302
	MQ2	NNR batchsize=4096	0.3508
	MQ3	NNC batchsize=4096	0.3556
	MQ4	BERT untrained	0.3359
	MQ5	PPO word2vec	0.3386
2	MQ1	First n_b	0.2897
	MQ2	NNR batchsize=4096	0.3360
	MQ3	NNC batchsize=4096	0.3376
	MQ4	BERT untrained	0.3030
	MQ5	PPO word2vec	0.2662
3	MQ1	First n	0.3506
	MQ2	NNR batchsize=4096	0.3729
	MQ3	NNC batchsize=4096	0.3651
	MQ4	BERT untrained	0.3506
	MQ5	PPO word2vec	0.3018
4	MQ1	PPO BERT	0.2975
	MQ2	NNR batchsize=4096	0.2973
	MQ3	NNC batchsize=4096	0.2987
	MQ4	BERT untrained	0.2954
	MQ5	PPO word2vec	0.2585
5	MQ1	PPO BERT	0.3135
	MQ2	NNR batchsize=4096	0.3237
	MQ3	NNC batchsize=4096	0.3276
	MQ4	BERT untrained	0.3316
	MQ5	PPO word2vec	0.3096

All runs in Table 6 have been described except for First n_b in batch 2. This is the same as First n in all other batches but the data has been pre-processed differently.

The classifier system (NNC) produces the best results in most runs, and this is consistent with our cross-validation results (Table 3). However, in contrast with Table 3, in most runs the second best is the regression system (NNR) instead of the untrained BERT system. All systems were re-trained for the BioASQ runs, keeping the same hyperparameters, so that we could use the entire training data, and it is possible that bad luck played a part here, and the BERT system was not trained to its best.

The reinforcement learning experiments show the lowest evaluation results. We should note, however, that the RL runs submitted to BioASQ7b also had lower results than the other runs, but the human evaluation results ranked them higher than our other runs. Also, the reinforcement learning submissions were not re-trained using the entire training data, which may be worth exploring because PPO had a more consistent learning curve for us than REINFORCE. We are waiting for the human evaluation results of BioASQ8b with anticipation.

Table 7 shows the results of the Australian National University’s submissions to BioASQ8b.

Table 7. SBERT runs submitted to BioASQ8b

Batch	Run Name	Description	ROUGE-SU4 F1
2	sbert reg	SBERT R 1 epoch	0.2856
	sbert cls	SBERT C 1 epoch	0.2958
	multitask sbert reg	SBERT M R 1 epoch	0.2857
	multitask sbert cls	SBERT M C 1 epoch	0.2910
3	sbert reg	SBERT R 20 epochs	0.3241
	sbert cls	SBERT C 20 epochs	0.3349
	multitask sbert reg	SBERT M R 20 epochs	0.3342
	multitask sbert cls	SBERT M C 20 epochs	0.3506
	sbert 1 epoch cls	SBERT C 1 epoch	0.3395
4	sbert reg	SBioBERT R 20 epochs	0.2601
	sbert cls	SBioBERT C 20 epochs	0.2601
	multitask sbert reg	SBioBERT M R 20 epochs	0.2572
	multitask sbert cls	SBioBERT M C 20 epochs	0.2527
	sbert cls 1 epoch	SBioBERT C 1 epoch	0.2580
5	sbert reg	SBioBERT R 20 epochs	0.3201
	sbert cls	SBioBERT C 20 epochs	0.3235
	multitask sbert reg	SBioBERT M R 20 epochs	0.3152
	multitask sbert cls	SBioBERT M C 20 epochs	0.3245
	sbert cls 1 epoch	SBioBERT C 1 epoch	0.3231

We observe that the SBERT runs perform worse or on par with the BERT untrained model of Table 6. This might be because questions and answers may not appear closely in an embedding space, and therefore an identical processing of each of them might not be advantageous. For example, a question includes question words whereas the candidate sentences are not normally questions. We also observe that the BioBERT models seem to produce lower results when compared with the BERT untrained model of Table 6. In general, the multitask learning approach performed better than just classification or regression.

8 Conclusions

This paper presents our approaches to query focused multi-document extractive summarisation for BioASQ8b. Our experiments include the use of BERT and BioBERT, Siamese architectures and SBERT, and Reinforcement Learning with PPO.

We observed that an approach that uses BERT to obtain the word embeddings, followed by LSTM to map these word embeddings to sentence embeddings, had the most promising results. The variant with BioBERT did not present an improvement, and this conflicts with the overall improvement of the use of BioBERT for question answering on biomedical texts. The approaches with Siamese architectures did not present an improvement over the base versions, presumably because questions and candidate sentences have different sentence styles.

As further research we plan to explore fine-tuning techniques for the BERT and BioBERT variants.

Acknowledgements

Research by Vincent Nguyen is supported by the Australian Research Training Program and the CSIRO Postgraduate Scholarship.

References

1. Dang, H.T.: Duc 2005: Evaluation of question-focused summarization systems. In: Proceedings of the Workshop on Task-Focused Summarization and Question Answering. pp. 48–55 (2006)
2. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: NAACL-HLT. pp. 4171–4186. Minneapolis, MN (jun 2019), <https://www.aclweb.org/anthology/N19-1423/>
3. Hill, A., Raffin, A., Ernestus, M., Gleave, A., Kanervisto, A., Traore, R., Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., Wu, Y.: Stable baselines. <https://github.com/hill-a/stable-baselines> (2018)

4. Howard, J., Ruder, S.: Universal language model fine-tuning for text classification. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 328–339. Melbourne, Australia (Jul 2018). <https://doi.org/10.18653/v1/P18-1031>, <https://www.aclweb.org/anthology/P18-1031>
5. Jones, C.R.: Reinforcement Learning For Query-based Multi-document Extractive Summarisation. Master’s thesis, Macquarie University (Jan 2020)
6. Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C.H., Kang, J.: BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* **36**(4), 1234–1240 (09 2019), <https://academic.oup.com/bioinformatics/article/36/4/1234/5566506>
7. Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., Zettlemoyer, L.: BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 7871–7880. Online (Jul 2020), <https://www.aclweb.org/anthology/2020.acl-main.703>
8. Liu, Y., Lapata, M.: Text summarization with pretrained encoders. 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). pp. 3730–3740. Hong Kong, China (Nov 2019). <https://doi.org/10.18653/v1/D19-1387>, <https://www.aclweb.org/anthology/D19-1387>
9. Mekala, D., Shang, J.: Contextualized weak supervision for text classification. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 323–333. Online (Jul 2020), <https://www.aclweb.org/anthology/2020.acl-main.30>
10. Mollá, D., Jones, C.: Classification betters regression in query-based multi-document summarisation techniques for question answering. In: Cellier, P., Driessens, K. (eds.) *Machine Learning and Knowledge Discovery in Databases*. pp. 624–635. Springer International Publishing, Cham (2020)
11. Reimers, N., Gurevych, I.: Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. 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). pp. 3980–3990. Stroudsburg, PA, USA (2019). <https://doi.org/10.18653/v1/D19-1410>, <https://www.aclweb.org/anthology/D19-1410>
12. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017)
13. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* **8**(3-4), 229–256 (1992)
14. Yoon, W., Lee, J., Kim, D., Jeong, M., Kang, J.: Pre-trained language model for biomedical question answering. In: Cellier, P., Driessens, K. (eds.) *Machine Learning and Knowledge Discovery in Databases*. pp. 727–740. Springer International Publishing, Cham (2020)