

# NLP-IIS@UT at SemEval-2021 Task 4: Machine Reading Comprehension using the Long Document Transformer

**Hossein Basafa**  
College of ECE  
University of Tehran  
hbasafa@ut.ac.ir

**Sajad Movahedi**  
College of ECE  
University of Tehran  
s.movahedi@ut.ac.ir

**Ali Ebrahimi**  
College of ECE  
University of Tehran  
ali96ebrahimi@ut.ac.ir

**Azadeh Shakery**  
College of ECE  
University of Tehran  
shakery@ut.ac.ir

**Heshaam Faili**  
College of ECE  
University of Tehran  
hfaili@ut.ac.ir

## Abstract

This paper presents a technical report of our submission to the 4th task of SemEval-2021, titled: Reading Comprehension of Abstract Meaning. In this task, we want to predict the correct answer based on a question given a context. Usually, contexts are very lengthy and require a large receptive field from the model. Thus, common contextualized language models like BERT miss fine representation and performance due to the limited capacity of the input tokens. To tackle this problem, we used the longformer model to better process the sequences. Furthermore, we utilized the method proposed in the longformer benchmark on wikipop dataset which improved the accuracy on our task data from (23.01% and 22.95%) achieved by the baselines for subtask 1 and 2, respectively, to (70.30% and 64.38%).

## 1 Introduction

Reading comprehension is the ability to understand a passage either by human or machine. One of the great benchmarks to evaluate this ability is to try to answer specific questions related to the passage (Rajpurkar et al., 2016). Generally, this problem can contain single or multiple documents as context (containing relevant information needed to understand and answer the question), a question (a sentence with at least one asking parameter), and an answer (which is the parameter value of the question).

In the Task of Reading Comprehension of Abstract Meaning (ReCAM), we have one passage as a context, one question and five candidate answers (Zheng et al., 2021). The goal is to identify the correct answer based on the context and the given question. You can see a sample of the data in Table 1. For each instance of the data, there is a passage, a question with a missing word that should be filled based on the passage, and five candidate answers to the question.

<b>Passage</b>	... observers have even named it after him, “Abenomics”. It is based on three key pillars - the “three arrows” of monetary policy, fiscal stimulus and structural reforms in order to ensure long-term sustainable growth in the world’s third-largest economy. In this weekend’s upper house elections ...
<b>Question</b>	Abenomics: The @Placeholder and the risks
<b>Answer</b>	(A) chances (B) prospective (C) security (D) objectives (E) threats

Table 1: An instance of the data.

The task divides into two subtasks: imperceptibility and non-specificity (Zheng et al., 2021).

- imperceptibility: this level of abstract words refers to ideas and concepts that are distant from immediate perception; such as culture, economics, and politics.
- non-specificity: In contrast to concrete words, this subtask includes more abstract words which focus on a different type of definition; for example, a concrete word like ‘cow’ could be interpreted as an ‘animal’ which is considered as a more abstract word (Changizi, 2008).

The main challenges of this task are the abstract meaning concept representation as well as the machine reading comprehension. This is the main reason we have utilized contextualized language representation models to tackle abstract meaning representation problems.

In this paper, we use an end-to-end deep contextualized architecture to model this task. This model is also capable of considering more than

one passage as the context, and more than five candidate answers. Since we use the long document transformer model (Longformer (Beltagy et al., 2020)), no limitation is considered in context passage length. We have evaluated this model both on subtask-1 and subtask-2 which resulted in 70% and 64% accuracy, respectively. Therefore, we have about 40% improvement compared to the baseline, which is a Gated Attention (GA) model (Zheng et al., 2021).

The rest of the paper is as follows: Section 2 describes the related works and the background. Section 3 includes the description of the proposed method. Section 4 contains the evaluation metrics used as well as a brief discussion, which is then followed by a conclusion and future works in section 5.

## 2 Background and Related Works

Many approaches have been presented in the literature, from pipeline-based models to end-to-end ones. Each module is also well-investigated from rule-based models to deep learning ones. Despite various configurations presented in the literature to model this problem, most of the systems consist of three modules (Baradaran et al., 2020):

- Language representation: this module is responsible to encode the inputs. Context, question, and answer need to be represented as numeric values for computational algorithms to be usable on them. Dense vectorized representations are the most popular methods, which allow us to use the majority of machine learning algorithms.
- Reasoning: this module is used to find demonstrations of why the answer is assumed to be valid. It can also be used as a limiter for searchable context.
- Prediction: this module aims to generate, retrieve or select the correct answer based on the task description.

Recent studies are provided as follows with respect to these modules that the last two modules have been merged. In the end, the longformer model is presented as our mainstay in this paper.

### 2.1 Word and text representation

One of the most important problems in NLP is representation learning. The earliest models for word

representation in the time of deep learning were the models proposed in (Pennington et al., 2014) and (Mikolov et al., 2013), which utilized the weights learned for an auxiliary task (a simplified version of the task of language modeling) for word representation. Similarly, methods proposed in (Le and Mikolov, 2014) and (Liu et al., 2015) utilized a similar structure for sentence, paragraph, or document representation learning.

While these methods were quite effective, it has been shown that using neural language models as a way of word representation results in much better, and context-aware representations. In (Howard and Ruder, 2018) it has been shown that fine-tuning language models as sentence encoders result in a significant performance improvement. At the same time, (Peters et al., 2018) used language models directly as word representations, which resulted in significant improvements. In (Devlin et al., 2018) a transformer model was trained for the task of masked language models, which resulted in significant improvements, surpassing human performance in many NLP tasks. One of the shortcomings of transformers is the lack of a memory mechanism, which results in (theoretically) lower receptive field compared with LSTMs (Beltagy et al., 2020) this shortcoming was addressed by improving the self attention mechanism in transformers so that it would have a (theoretically) unbounded receptive field. More details are presented later in this section.

### 2.2 Natural language understanding

Natural language understanding (NLU) is an umbrella term, referring to any tasks that require machine comprehension. Compared to other NLP tasks, NLU requires the model to be able to understand and reason about the data (Semaan, 2012). While great progress has been made in this field by using contextual word representation (Devlin et al., 2018), it has been found that designing the model itself must not be neglected (Zhu et al., 2018). On the other hand, it has been shown that utilizing a transfer learning setting to share knowledge between different NLU tasks results in better performance with fewer data and fewer parameters (Pilault et al., 2020), which proves a significant similarity between these tasks.

### 2.3 The Longformer

Deep contextualized language models like BERT (Devlin et al., 2019) have been well investi-

gated in the literature and achieved state-of-the-art results on various tasks. However, these models suffer from performance limitations due to their self-attention layer which results in quadratic space and time complexity concerning the sequence length. In contrast, this model removes the self-attention layer from the base language models, so the limitation resolves and the complexity scales to linear. In order to increase the quality of the model compared to basic models, they have added a global attention layer to the model end which significantly outperforms state-of-the-art models on long document (passage) tasks and competitive on normal documents. Also, this configuration increases the performance on both normal and lengthy inputs which makes it a good alternative for tasks containing large inputs. This model is also evaluated on a similar task on WikiHop dataset (Welbl et al., 2018) and improved the results in terms of accuracy (Beltagy et al., 2020).

### 3 Method

As mentioned in section 1, given a passage, a question, and a set of answers to the question, the goal is to predict the correct answer among the candidates, which can be seen as a benchmark to evaluate how well the model can comprehend the abstract meaning. To do so, we considered an end-to-end deep learning architecture based on the transformer architecture.

Specifically, we used contextual word embeddings based on the transformer to better discover and encode the information contained in the passage. In our model, both subtasks use the same architecture as shown in figure 1, although we did not experiment on the possibility of multi-task learning. The word representation models are fine-tuned on the data for better performance. The fine-tuning procedure could allow us to extract additional, task-related information which could result in better accuracy in the evaluation phase.

To model this problem, let  $c = \{c_1, c_2, \dots, c_I\}$  denote the passage as the context, where  $c_i$  corresponds to the  $i^{th}$  token (word or subword, depending on the tokenization technique used) and  $I$  is the number of tokens in the passage. Similarly, the question is considered as  $q = \{q_1, q_2, \dots, q_K\}$  where  $K$  denotes the length of the question, and  $q_k$  corresponds to the  $k^{th}$  token of the question. Each answer also denotes as  $e^j$  which is only one abstract word ( $j \in \{1, 2, \dots, 5\}$ ). Then we concatenate

the question and the candidates as:

$$a = [q; e^1; e^2; \dots; e^5]. \quad (1)$$

The size of this sequence is  $A = K + 5$  as we have only 5 candidates. Generally, this can be an arbitrary length based on the dataset.

Note that we introduce special tokens to separate the context, the question, and the candidates, similar to (Beltagy et al., 2020). Specifically, we introduce the tokens  $\langle s \rangle$  and  $\langle /s \rangle$  for separating the context,  $\langle q \rangle$  and  $\langle /q \rangle$  for separating the the question, and the tokens  $\langle ent \rangle$  and  $\langle /ent \rangle$  for separating the candidates from each other. In the case of multiple passages, all passages are concatenated to form a single context. These tokens are randomly initialized and fine-tuned.

We used the Longformer model introduced in (Beltagy et al., 2020) as the pre-trained contextual embedding model in our method. Since the context could be too long, we split the context sequence to separate chunks. Each chunk length is equal to maximum sequence length the model could accept appending the sequence  $a$ ; in fact,  $model\_max\_length = len(chunk) + len(a)$ . If  $c^l$  denote each chunk, this sequence could be showed as:

$$b = [c^l; a] \quad (2)$$

where the full context is  $c = \{c^1, c^2, \dots, c^L\}$ , and  $L$  is the last chunk. The size of this sequence is  $B$  so  $B = L + A$ .

After feeding the input  $b$  to the Longformer model, we apply a global attention only on  $a$  (concatenated question and answer candidates), and the rest is the context. As the longformer model utilizes a base model (like RoBERTa without the self-attention layer, in our case), we denote this as *basemodel* function that outputs the encoded sequence of the input. If  $GAttn$  denotes the global attention function, we have:

$$d_i = basemodel(b) \quad (3)$$

$$g_i = GAttn(d_i).1(i \in A) \quad (4)$$

where  $d_i$  is the raw output vector for each input token. The global attention function is applied if it is a question or answer candidate token. Then, we extract the outputs corresponding to the question and the candidates tokens, i.e. we have:

$$h_j = GAttn(a, c^l) \quad (5)$$

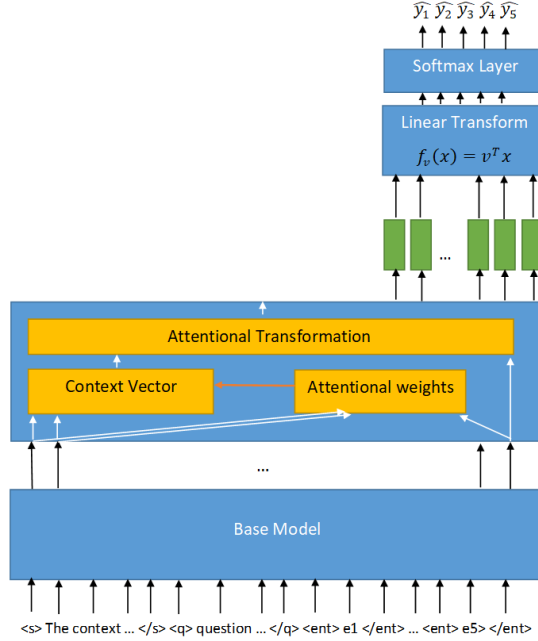


Figure 1: The model architecture. The concatenated input vector will be encoded using the base model (like RoBERTa without the self-attention layer, in our case). A global attention(Luong et al., 2015) will be applied to the question and the candidate answers representations with respect to the passage as the context. The logit (score) of each ent token will be calculated using a linear transformation function, then the prediction distribution over the answer candidates (ent tokens) will be outputted using a softmax layer.

Finally, we obtain the logit of each candidate (<ent> tokens) as  $x_j$  ( $x_j = h_j$  if  $j$  correspond to a candidate), average over different chunks, and apply a linear transformation:

$$f_j = v^T x_j \quad (6)$$

where the vector  $v$  is trainable, and  $f_j$  is the score of each candidate. And the probability distribution over the candidates will be calculated using a softmax layer on the logits. The predicted answer is the argmax of the softmax output. we fine-tuned the model using the cross-entropy loss.

## 4 Evaluation

Although we only participated in the second subtask, we will evaluate our model on both subtasks here. We will explain our configurations for utilizing the model on the task as well as other baselines which are the BERT-base as an alternative model and the Gate-Attention (GA) as our task baseline. Finally, a brief discussion will be done based on the results.

### 4.1 Metrics

Popular metrics to evaluate these models are F1, EM (Exact Match or accuracy), and MRR (Mean

Reciprocal Rank). As the precision and recall in our task are equal, so  $F1 = \text{Precision} = \text{Recall}$ . Also, F1 and EM are the same. And, the use of MRR is optional, so the metrics used to evaluate the result are the accuracy and the F1.

### 4.2 Baseline configuration

The baseline model (GA) is trained for 30 epochs, each epoch containing 101 mini-batches. The train batch size is set to 32. Dropout with the rate of 0.5 is also applied to the hidden states, and the learning rate is set to 0.001. The dimensionality of the GloVe embedding is 300, and the hidden size is set to 128. Training and evaluation take about 2 hours on a single v100 GPU.

### 4.3 BERT configuration

We use the same configuration as our method except for the global attention mechanism. In fact, we consider the output vector of each chunk as our final vector to be linearly transformed into single logit, followed by a softmax layer using the cross-entropy loss. Similarly, the logit is averaged over different chunks, before applying the linear transformation. Note that the maximum sequence length here is bounded to 512 tokens, and the model includes the  $n^2$  attention mechanism. We use the

Metrics	Baseline(GA)	BERT	Our Method
Accuracy	23.01%	63.43%	<b>70.30%</b>
Macro Avg F1	22.83%	63.38%	<b>70.23%</b>
Weighted Avg F1	22.76%	63.40%	<b>70.27%</b>

Table 2: Subtask1 evaluation metrics on the test set

Metrics	Baseline(GA)	BERT	Our Method
Accuracy	22.95%	58.76%	<b>64.38%</b>
Macro Avg F1	22.42%	58.72%	<b>64.35%</b>
Weighted Avg F1	22.45%	58.75%	<b>64.40%</b>

Table 3: Subtask-2 evaluation metrics on the test set

base version of the model and fine-tuned it on each subtask.

#### 4.4 Our method configuration

We used the same model introduced in section 3 for both subtasks. The model was initialized using the Longformer-base pre-training weights, then fine-tuned in each of the subtasks. Due to the performance issues, the model max sequence length is set to 4096 tokens which are sufficient in our case. We also used the RoBERTa-large tokenizer to tokenize the input sequence as the Longformer model has been trained on using this configuration. We used a batch size of 32 and a maximum learning rate of  $3e-5$  using the Adam optimizer with  $\beta_2=0.98$ . We then assumed the validation check interval to 250 which indicates the number of gradient updates between checking validation loss. And a weight decay of 0.01 has been considered to regularize the model and avoid overfitting.

Our proposed model is trained for 15 epochs for each task. Fine-tuning the model takes about six hours, and inference takes about nine seconds for each sample on a single V100 GPU.

#### 4.5 Evaluation of Subtask 1

Subtask1 measures imperceptibility abstract level of language understanding. This subtask includes 3227 training samples, 837 validation samples, and 2025 test samples. The size of the biggest sample in terms of context length is about 2000 tokens. We have achieved an accuracy of 70% on the validation set, which improves our baseline by about 40 percent. Table 2 showed the results of this subtask.

#### 4.6 Evaluation on Subtask 2

Subtask2 measures the non-specificity level of abstract meaning in reading comprehension. It in-

cludes 3318 training samples, 851 validation samples, and 2017 test samples. The best accuracy on the validation set is 64%. Table 3 showed the results of this subtask.

#### 4.7 Discussion

We used two baselines to find out the effect of using a pre-trained model rather than a simple RNN model. Although this task offers a higher level of representation, using the pre-train models is helpful, and there is a higher chance of modeling such abstract concepts.

The results on subtask2 are weaker than subtask1 in pre-trained models. This can be the consequence of limited semantic representation for abstract word which indicates the subtask2 includes more abstract words in terms of abstract level; for example, the word 'animal' could be matched to any animal, like 'cat' or 'dog', but the word 'entity' is hard to be represented as it could be matched to a large number of words. And the model faces a limitation in the knowledge representation. Another assumption could be the data enrichment that these model has been trained on. As most of the available texts for training consist of concrete words, it is more likely to leverage the language understanding to less abstract words to achieve a better result.

Comparing our method which is based on longformer model to usual language models like BERT indicates a new insight in terms of passage length and the attention mechanism. Popular language models like BERT and RoBERTa use a  $n^2$  attention which requires a large receptive field to represent long passages. This results in the performance limitation which bounds the input sequence up to 512 tokens. In contrast, the longformer global attention mechanism relaxes this limitation as we only need to pay attention to a small factor of context

and more focus on the local window. So the receptive field will not overflow and saves the necessary information to better represent the language.

We have analyzed the errors that mostly affect our model performance. We think that the problem is the contextual representation of the language modeling, which is not well-suited in our method i.e. concatenating the context, question, and answer. The main disadvantage of concatenating the candidate answers to each other is the missing fine contextual representation as the state-of-the-art models consume the position embedding. Additionally, incorrect candidates register additional noise to each word representation as well as the placeholder in the question.

## 5 Conclusion and Future works

We have shown how different approaches can be leveraged to machine reading comprehension of abstract meaning. We reformulated the longformer model to learn abstract meaning as a new level of semantic in machine reading comprehension. This method can also be improved by taking advantage of external knowledge and task-specific model architectures that optimize the current baseline.

## References

- Razieh Baradaran, Razieh Ghiasi, and Hossein Amirkhani. 2020. [A survey on machine reading comprehension systems](#). *CoRR*, abs/2001.01582.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Mark A Changizi. 2008. Economically organized hierarchies in wordnet and the oxford english dictionary. *Cognitive Systems Research*, 9(3):214–228.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196. PMLR.
- Yang Liu, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. 2015. Topical word embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. *arXiv preprint arXiv:1310.4546*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Jonathan Pilault, Amine Elhattami, and Christopher Pal. 2020. Conditionally adaptive multi-task learning: Improving transfer learning in nlp using fewer parameters & less data. *arXiv preprint arXiv:2009.09139*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100, 000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392. The Association for Computational Linguistics.
- Paul Semaan. 2012. Natural language generation: an overview. *J Comput Sci Res*, 1(3):50–57.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. Constructing datasets for multi-hop reading comprehension across documents. *Transactions of the Association for Computational Linguistics*, 6:287–302.
- Boyuan Zheng, Xiaoyu Yang, Yuping Ruan, Quan Liu, Zhen-Hua Ling, Si Wei, and Xiaodan Zhu. 2021. SemEval-2021 task 4: Reading comprehension of abstract meaning. In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*.
- Chenguang Zhu, Michael Zeng, and Xuedong Huang. 2018. Sdnet: Contextualized attention-based deep network for conversational question answering. *arXiv preprint arXiv:1812.03593*.