

CS-NET at SemEval-2020 Task 4: Siamese BERT for ComVE

Soumya Ranjan Dash* **Sandeep Routray*** **Prateek Varshney*** **Ashutosh Modi**

Indian Institute of Technology Kanpur (IITK)

{soumyard, sroutray}@iitk.ac.in, varshney@cse.iitk.ac.in
ashutoshm@cse.iitk.ac.in

Abstract

In this paper, we describe our system for Task 4 of SemEval 2020, which involves differentiating between natural language statements that confirm to common sense and those that do not. The organizers propose three subtasks - first, selecting between two sentences, the one which is against common sense. Second, identifying the most crucial reason why a statement does not make sense. Third, generating novel reasons for explaining the against common sense statement. Out of the three subtasks, this paper reports the system description of subtask A and subtask B. This paper proposes a model based on transformer neural network architecture for addressing the subtasks. The novelty in work lies in the architecture design, which handles the logical implication of contradicting statements and simultaneous information extraction from both sentences. We use a parallel instance of transformers, which is responsible for a boost in the performance. We achieved an accuracy of 94.8% in subtask A and 89% in subtask B on the test set.

1 Introduction

Incorporating common sense in natural language understanding systems and evaluating whether a system has sense-making capability remains a fundamental question in the natural language processing field (Modi, 2017; Modi, 2016; Modi and Titov, 2014). One important difference between human and machine text understanding lies in the fact that humans have access to commonsense knowledge while processing text, which helps them to draw inferences about facts that are not mentioned in a text, but that is assumed to be common ground (Modi et al., 2017). For a computer system, inferring unmentioned facts is a non-trivial challenge (Ostermann et al., 2018a). For our problem, we have proposed methods to include common sense in the validation and reasoning paradigm (Wang et al., 2019).

Task 4 of semeval 2020 (Wang et al., 2020) is a common-sense validation and explanation task. It consists of classifying against common sense sentences from sentences that make sense. Figure 1 shows examples from subtask A and subtask B. In subtask A, clearly sentence 1 is against common sense. Subtask B contains three options for reasons to explain why sentence 1 is against common sense. As orange juice does not taste good on cereal, but milk does, sentence 1 makes less sense than sentence 2.

We use the generated embedding from transformer based encoders like BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), AIBERT (Lan et al., 2019) which capture in-context semantic information and design a siamese architecture to extract the relational information between the sentences. The implementation for our system is made available via Github¹.

2 Background

2.1 Problem Definition

The task is to develop a system that can differentiate natural language statements that make sense from those that do not. The organizers present three subtasks to test these. Out of these, we describe the

* Authors equally contributed to this work.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

¹<https://github.com/soumyardash/SemEval2020-Task4>

Subtask A	Subtask B
Sentence 1: He poured orange juice on his cereal. Sentence 2: He poured milk on his cereal. ✓	Sentence: He poured orange juice on his cereal. Reason 1: Orange juice is usually bright orange. Reason 2: Orange juice doesn't taste good on cereal. ✓ Reason 3: Orange juice is sticky if you spill it on the table.

Figure 1: Sample data point explaining the problem

description of the system proposed for the first two subtasks, subtask A and subtask B. Each instance in the given dataset is composed of 10 sentences: $\{S_1, S_2, O_1, O_2, O_3, R_1, R_2, R_3\}$. S_1 and S_2 are two syntactically similar statements, differing by only a few words, with only one making sense.

- **Subtask A Validation:** Requires the model to choose which of the two statements S_1 and S_2 does not makes sense. We frame it as a binary classification problem and estimate the probability that the sentence is against common sense. Accuracy score is used for evaluation.
- **Subtask B Explanation (Multi-Choice):** Requires the model to choose the most appropriate of the three reasons $\{O_1, O_2, O_3\}$ to explain the against-common-sense statement (one of S_1 or S_2). We formulate this as a multi-class classification problem and estimate the probability that the reason is in fact the correct explanation. Accuracy score is used for evaluation.

2.2 Related Work

Language models : Over the years, a great amount of effort has been directed towards creating benchmark datasets that can measure a system’s performance on language processing tasks and provide an impetus for the development of new approaches to the tasks. These benchmark tasks have led to many computational models ranging from earlier symbolic and statistical approaches to recent approaches based on deep neural networks; which model context of language, take advantage of external data or knowledge resources, and achieve the state-of-the-art performance, and at times, even near or above human performance. A major landmark in NLP is the development of pre-trained models and embeddings that can be used as features or further fine-tuned for downstream tasks. These models are often trained based on large corpora of textual data to capture different word senses.

The defining contribution of Embeddings from Language Models (ELMo) (Peters et al., 2018) is its contextual word embeddings, which are built relying on the entire input sentence that they belong to. The recent Bidirectional Encoder Representations from Transformers (BERT) model outperforms previous competitive approaches by better capturing the context. A more recent model, XLNET (Yang et al., 2019), exceeded the performance of the vanilla BERT variant on several benchmarks. Robustly Optimized BERT Approach (RoBERTa) (Liu et al., 2019) achieved further improvement by making changes to the pre-training approach used in BERT. It includes randomizing masked tokens in the cloze pre-training task for each epoch instead of keeping them the same over epochs. It also augments the next sentence prediction pre-training task with an additional task which compels the model to also predict whether a candidate next sentence comes from the same document or not. A Lite BERT (ALBERT) (Lan et al., 2019) implements several novel parameter reduction techniques to increase the training speed and efficiency of BERT, enabling a much deeper scale-up than the original large variant of BERT while having fewer parameters. ELECTRA (Clark et al., 2020) is used to pre-train transformers with comparatively less computation. This model is similar to the discriminator of a GAN.

Common sense validation : TriAN (Wang et al., 2018) achieved state-of-the-art performance for SemEval ’18 Task-11 : Machine Comprehension Using Commonsense Knowledge (Ostermann et al., 2018b). It proposed a threeway attention mechanism to model interactions between the text, question, and answers, on top of BiLSTMs. It incorporated relational features (based on ConceptNet).

CommonsenseQA (Talmor et al., 2018) is a large multi-choice question answering dataset that requires different types of commonsense knowledge to predict the correct answers. The leaderboard provides a variety of ensembling techniques using different LMs and also several instances of the same LM.

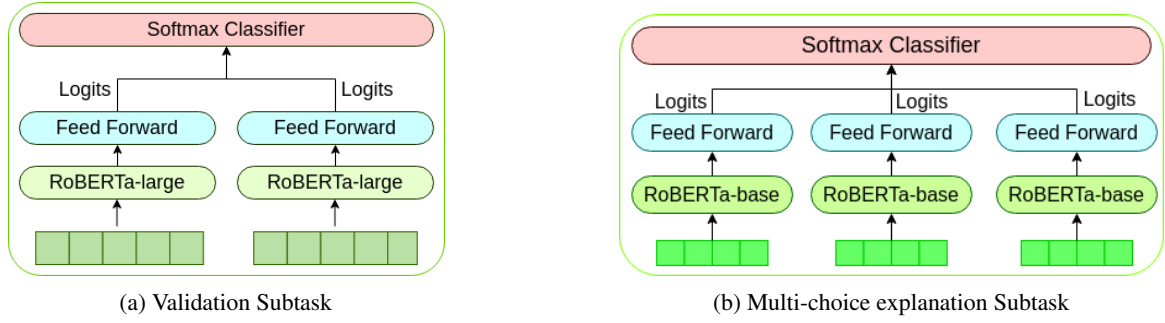


Figure 2: Final architecture used for the subtasks

3 System Overview

We mention some crucial insights from our initial experimentation that were helpful to arrive at the final approach. We consider both subtasks separately.

3.1 Subtask A: Validation

3.1.1 Initial Experimentation

Initially, we experimented with the vanilla implementation of BERT sequence classifier model. It consists of a linear classification head over the BERT embedding of the [CLS] token which are supposed to have learnt representation for the entire sentence. The classifier predicts a binary label indicating if the input sentence made sense or not. The weights of the BERT encoder stack are fixed in this approach. Even though the accuracy obtained (see Table 2a) was higher than the baseline, it was still on the lower side.

Improving upon the initial approach, we then updated the weights of BERT encoder during training phase. We also modified the training samples as - if S_1 and S_2 are the given sentence pair, we concatenated these with the help of a phrase, e.g. S_1 makes more sense than S_2 , and S_2 makes more sense than S_1 . The task was now to predict binary label depending on the correctness of the claim. These changes resulted in a significant gain in the performance. Interestingly, we observed that the performance varied slightly depending on the conjunctive phrase used to join the sentences. While conjunctive phrases like "makes more sense than" gave better results, the phrases like "because" and "not valid" had a lower performance. Further, we concluded that training the BERT encoder was necessary to achieve good performance.

We also observed that the approach proposed above suffered from logical fallacy, that is, in case of mispredictions the model assigned both sentences of the pair as against commonsense. Further, though the approach worked well in the trainset, it did not generalize well to the devset, indicating overfitting.

3.1.2 Proposed Approach

To address the issues discussed above, we designed a siamese architecture (see Fig. 2a). The model took both the sentence at the same time for prediction and predicted a probability score over the two sentences. Mathematically,

$$x_1 = E(s_1) \quad x_2 = E(s_2) \quad l_1 = \frac{e^{Wx_1}}{e^{Wx_1} + e^{Wx_2}} \quad l_2 = \frac{e^{Wx_2}}{e^{Wx_1} + e^{Wx_2}}$$

where s_1 and s_2 are vectorized representations of tokens of S_1 and S_2 . $E(\cdot)$ denotes the trainable transformer encoder block weights and W is the trainable classifier weights. x_1 and x_2 are the representations of the starting [CLS] token. Cross-entropy loss was used while training to determine the optimal weights.

Note that $E(\cdot)$ and W are same for both the sentences S_1 and S_2 . Since, the weights of the two networks (transformer based encoder+classifier) were shared, the model was able to focus better on the features which differed between the two sentences and avoid overfitting. Such features proved crucial in deciding the output label which was supported by the superior performance on both trainset and devset. We also experimented with variants of BERT: BERT-base, BERT-large, RoBERTa-base and RoBERTa-large. The best performance was obtained with RoBERTa-large, which formed our final submission. We did not observe any significant gains in using conjunctive phrases and hence excluded it from our final submission.

3.2 Subtask B: Multi-Choice

3.2.1 Initial Experimentation

As an initial approach, we formulated the task as a binary classification problem using the BERT sequence classifier model. Using the given against commonsense sentence S from the dataset and the three candidate reasons $O_i, i = 1, 2, 3$, we formed sentences S makes less sense because O_i . Although the performance obtained using this approach was above the baseline, it was still on the lower side (see Table2b).

Next, we tried using conjunctive phrases. We formed *The reason S makes less sense is O_1 rather than O_2 or O_3* , with all possible permutations of O_1, O_2, O_3 and used it as our input instead of original sentences. The idea was to allow the model to access information from other candidate reasons while deciding the correctness of a particular reason. This approach gave a good performance on the trainset but did not perform well on the devset.

Above approaches also suffered from similar logical fallacy, as described in Section 3.1.1. Motivated by the success of siamese architecture in avoiding over-fitting, enabling efficient information sharing, and eliminating logical fallacies, we designed a three-way siamese architecture for subtask B (See Fig.2b).

3.2.2 Proposed Approach

First, we formed the three inputs by joining S with O_1, O_2, O_3 respectively using a separator token ([SEP]). We did not use any conjunctive phrases in our final approach. The model took the inputs at the same time and predicted a probability score indicating the likelihood of the reason being the correct explanation. Mathematically,

$$\begin{aligned} x_1 &= E(s_1) & x_2 &= E(s_2) & x_3 &= E(s_3) \\ l_1 &= \frac{e^{Wx_1}}{\Sigma} & l_2 &= \frac{e^{Wx_2}}{\Sigma} & l_3 &= \frac{e^{Wx_3}}{\Sigma} \end{aligned}$$

where s_1, s_2, s_3 are the vectorized representation of input tokens and $\Sigma = e^{Wx_1} + e^{Wx_2} + e^{Wx_3}$. $E(\cdot)$ denotes the trainable transformer encoder block weights and W is the trainable classifier weights. Cross-entropy loss function was used while training to determine the optimal weights.

The shared weight among the inputs enables efficient information exchange. The model actually compares the reasons against each other while determining the correct explanation. We experimented with using the [CLS] token representation and an average pooled representation of all the tokens as the input to feed-forward classification layer. We observed that better performance was obtained with an average pooling. The performance gains can be attributed to the fact that average pooling preserved more useful information for the classification layer, which was otherwise lost. We also experimented with variants of BERT: BERT-base, BERT-base, ALBERT-base, and RoBERTa-base. The best performance was obtained with RoBERTa-base and average pooling, which formed our final submission.

4 Experimental Setup

Data: Dataset that was used to build the models were provided by the organizers as a part of the pilot study (Cunxiang, 2019). As a single datapoint, Subtask A had an against-common-sense and a correct sentence. For each of the against common sense sentence of Subtask A, Subtask B had a sensible reason and two confusing reasons. The datasets were provided in three phases. The training data was used to train the language model. It consisted of 10,000 datapoints. The 997 datapoints long development dataset was used to tune the hyperparameters of the language model. The answers generated on the test dataset of 1,000 datapoints were used for submission. From Table 1, we observe that the label distribution is unbiased for training as well as the development dataset.

Parameter setting: We used held-out validation using the development dataset for validation. The train:dev ratio for the training was nearly 10:1. The hyperparameters were tuned using a grid search around the default setting of the language model. It ensures the highest efficiency of the model. For both the subtasks, the accuracy of the model on the development dataset was used as a measure of performance. Our submitted model is trained on batch mode with a batch size of 32 using AdamW optimizer. The learning rate of the optimizer was set to be $2e-5$, and the adam epsilon was set to be $1e-8$.

Dataset	First sentence invalid	Second sentence invalid	Dataset	First reason correct	Second reason correct	Third reason correct
Training Data	4979	5021	Training Data	3195	3362	3443
Validation Data	518	479	Validation Data	344	327	336

(a) Subtask A dataset details

(b) Subtask B dataset details

Table 1: Dataset analysis

Model	Accuracy
BERT Classifier	77.1%
BERT Classifier + phrase concat.	84.3%
Albert-base Siamese	87.6%
BERT-base Siamese	88.6%
RoBERTa-base Siamese	90.7%
Electra-base Siamese	93.6%
RoBERTa-large Siamese	95.2%

(a) Subtask A accuracy on dev set

Model	Accuracy
BERT Classifier	77.3%
BERT Classifier + phrase concat	83.2%
BERT-base Siamese	84.3%
ALBERT-base Siamese	85.7%
RoBERTa-base Siamese	87.5%
Electra-base Siamese	87.7%
RoBERTa-base Siamese+avg. pool	89.7%

(b) Subtask B accuracy on dev set

Table 2: Results analysis of various models

5 Results

5.1 Result Analysis

Table 2 compiles the system performance on the development dataset for various approaches.

For subtask A, Table 2a demonstrates the effect of adding conjunctive phrases to the sentence pair over the vanilla implementation of BERT. It exhibits the gain in accuracy obtained from addressing the logical fallacy problem by using a two-way siamese network, as discussed previously. It compares the performance of various models, notable among which is Electra and RoBERTa, which outperform BERT.

Table 2b exhibits a similar trend for subtask B, with a gain in accuracy after the use of concatenating phrases with the sentence and all its reasons. The use of the three-way siamese network further improves performance. The table also shows the effect of different language models on the task when used with the siamese architecture. The addition of an average pooling layer over the language model embedding gave improved results in comparison to the use of embedding of the [CLS] token only.

5.2 Error Analysis

For subtasks A and B, we found our system to be performing well on leader-board. Our system ranked 10th in subtask A and 12th in subtask B. While our system gave an accuracy of 94.8% on test set in subtask A, the top performance was 97%. In subtask B, the accuracy of our system was 89%, the top performance was 95%. There was a great diversity of contexts in our dataset. Our model learnt representations to capture the context. We believe the performance could be further improved by directly considering the relationship established between the keywords in the sentences. One way to do this would be to encode word relationships in a sentence using a knowledge base like ConceptNet. Then an attention layer could be used to learn a joint representation using representations from BERT-variant and ConceptNet.

6 Conclusion

In this paper, we have presented our systems for the Commonsense Validation and Explanation Challenge in SemEval2020. Our approach for subtask A and subtask B achieves close to state-of-the-art results. Our approach demonstrates the advantage of using parallel instances of the transformer in terms of a performance gain in classification based tasks.

References

Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators.

- Wang Cunxiang. 2019. Semeval2020-task4-commonsense-validation-and-explanation. <https://tinyurl.com/wjaccm3>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach.
- Ashutosh Modi and Ivan Titov. 2014. Inducing neural models of script knowledge. In *Proceedings of the Conference on Computational Natural Language Learning (CoNLL)*, Baltimore, MD, USA.
- Ashutosh Modi, Ivan Titov, Vera Demberg, Asad Sayeed, and Manfred Pinkal. 2017. Modelling semantic expectation: Using script knowledge for referent prediction. *Transactions of the Association for Computational Linguistics*, 5:31–44.
- Ashutosh Modi. 2016. Event embeddings for semantic script modeling. In *Proceedings of the Conference on Computational Natural Language Learning (CoNLL)*, Berlin, Germany.
- Ashutosh Modi. 2017. *Modeling common sense knowledge via scripts*. Ph.D. thesis, Universität des Saarlandes.
- Simon Ostermann, Ashutosh Modi, Michael Roth, Stefan Thater, and Manfred Pinkal. 2018a. McScript: A novel dataset for assessing machine comprehension using script knowledge. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, Miyazaki, Japan.
- Simon Ostermann, Michael Roth, Ashutosh Modi, Stefan Thater, and Manfred Pinkal. 2018b. SemEval-2018 task 11: Machine comprehension using commonsense knowledge. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 747–757, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2018. Commonsenseqa: A question answering challenge targeting commonsense knowledge.
- Liang Wang, Meng Sun, Wei Zhao, Kewei Shen, and Jingming Liu. 2018. Yuanfudao at semeval-2018 task 11: Three-way attention and relational knowledge for commonsense machine comprehension.
- Cunxiang Wang, Shuailong Liang, Yue Zhang, Xiaonan Li, and Tian Gao. 2019. Does it make sense? and why? a pilot study for sense making and explanation.
- Cunxiang Wang, Shuailong Liang, Yili Jin, Yilong Wang, Xiaodan Zhu, and Yue Zhang. 2020. SemEval-2020 task 4: Commonsense validation and explanation. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding.