

# BEIKE NLP at SemEval-2022 Task 4: Prompt-Based Paragraph Classification for Patronizing and Condescending Language Detection

Yong Deng, Chenxiao Dou\*, Liangyu Chen  
Deqiang Miao, Xianghui Sun, Baochang Ma, Xiangang Li

AI Center@Beike, China

{dengyong013, douchenxiao001, chenliangyu003}@ke.com

## Abstract

PCL detection task is aimed at identifying and categorizing language that is patronizing or condescending towards vulnerable communities in the general media. Compared to other NLP tasks of paragraph classification, the negative language presented in the PCL detection task is usually more implicit and subtle to be recognized, making the performance of common text-classification approaches disappointed. Targeting the PCL detection problem in SemEval-2022 Task 4, in this paper, we give an introduction to our team’s solution, which exploits the power of prompt-based learning on paragraph classification. We reformulate the task as an appropriate cloze prompt and use pre-trained Masked Language Models to fill the cloze slot. For the two subtasks, binary classification and multi-label classification, DeBERTa model is adopted and fine-tuned to predict masked label words of task-specific prompts. On the evaluation dataset, for binary classification, our approach achieves an F1-score of 0.6406; for multi-label classification, our approach achieves an macro-F1-score of 0.4689 and ranks first in the leaderboard.

## 1 Introduction

Patronizing and Condescending Language (PCL) towards vulnerable communities in the media has become a hot issue, which is heatedly discussed in society at present. The language often refers to the discourse which is published with a mixture of pity and superiority towards unprivileged groups. Such attitude is always adopted unconsciously or even out of kindness, however, its negative effect indeed exacerbates social prejudice and routinizes discrimination towards disadvantaged people (Ng, 2007). Despite the PCL emerging in numerous media, its style can not be precisely captured by far, since the usage of PCL is commonly unintended,

unaffected, and subjective. This poses many challenges to the detection of PCL, which attracts great attention from the NLP community. Even though there exist substantial NLP studies on paragraph identification in various fields, research specifically for PCL identification has not been yet seriously introduced (Pérez-Almendros et al., 2020). Nowadays with social discrimination continuously rising, an effective approach, which can automatically identify PCL towards vulnerable communities, becomes more and more necessary and important to our society.

To encourage more research on the PCL problem, SemEval-2022 Task 4 (Pérez-Almendros et al., 2022) provides an English PCL dataset for language-modeling study and evaluation. The main task contains two PCL-classification subtasks, one for binary classification (SubTask 1) and the other for multi-label classification (SubTask 2). Given a paragraph, SubTask 1 is aimed to identify whether or not it is written with PCL style. F1 over the positive class is taken as the evaluation metric of this task. For the same paragraph, the goal of SubTask 2 is to classify which specific PCL categories it belongs to. In detail, there are 7 different PCL types needed to be recognized, and one paragraph can have up to 7 PCL labels at the same time. For evaluation, SubTask 2 introduces macro F1 over the 7 classes as the metric.

In this paper, we propose our approach targeting PCL detection, which is submitted to SemEval-2022 Task 4. To tackle the sparsity and implicitness issues of PCL language, we first reformulate the task as a specific form of cloze prompt, and then apply prompt-based learning on it to predict the appropriate label words and corresponding task labels. To take advantage of the pre-trained Language Model (LM), DeBERTa (He et al., 2020) is adopted as our based model, which improves BERT (Devlin et al., 2018) by introducing two novel techniques: disentangled attention and decoding enhanced masking.

\* Corresponding Author

The same classification method and model architecture are used for both subtasks to train and predict. Experiments are conducted to show the effectiveness of our approach. At last, during the evaluation phase, our approach achieves an F1 score of 0.6406 on SubTask 1 (rank 4th on the leaderboard), and a macro-F1 score of 0.4689 on SubTask 2 (rank 1st on the leaderboard).

## 2 Related Work

Condescending and patronizing treatment is a controversial social problem, which attracts attention from researchers of various fields (Margić, 2017; Huckin, 2002). To the NLP community, although extensive work on detecting different kinds of harmful language is presented, the targeted language styles are often explicit, aggressive, and flagrant, which can be perceived obviously. Unlike the previously studied language, the style of PCL language can not be easily sensed, since the usage of PCL is commonly unintended, unaffected, and subjective, which makes it a challenging identification problem. In recent years, more and more researchers (Mendelsohn et al., 2020; Sap et al., 2019) start researching this emergent NLP topic, but the room for progress is still large.

Paragraph classification is a basic and important NLP task, which has been continuously studied by industry and academia. Applications based on paragraph classification are increasingly permeating our lives, such as spam filtering (Kumar, 2020) and text sentiment analysis (Gao et al., 2019). Traditional classification models, such as SVM (Joachims, 1998) and XGBoost (Chen and Guestrin, 2016), predict the category of paragraph mainly based on the statistical features extracted from raw text data. With the appearance of BERT (Devlin et al., 2018), many researchers (Croce et al., 2020; Jin et al., 2020) achieve success on paragraph classification, by utilizing the contextualized word vectors of pre-trained language models.

In recent years, with the development of prompt-based learning (Schick and Schütze, 2020b,a; Brown et al., 2020), researchers start reformulating the paragraph classification problem as a masked language modeling problem. Unlike the traditional models, prompt-based models first predict the label word of mask in a prompt and then map the label word to the label of category by a verbalizer. Thanks to the power of pre-trained language models, prompt-based models (Hu et al., 2021; Gu et al.,

2021) demonstrate their strength in performing few-shot or even zero-shot learning on the scenarios of paragraph classification with few or no labeled data. The above studies have shown that prompt-based learning is highly effective to solve the problem of paragraph classification. To study the effect of prompt-based learning on identifying paragraphs of PCL, in the rest of this paper, our modeling method and experimental analysis are introduced.

## 3 Our Approach

In this section, we present our approach to utilize prompt-based learning for PCL detection. We first give the overall paradigm of our model and then introduce the ensemble strategy used in this task.

### 3.1 Prompt-Based Model

Unlike other traditional classification methods which take original text as input, the prompt-based classification method first wraps the original text with some specific-designed template and then feeds the synthesized text into a language model for classification. In this way, the paragraph classification problem is reformulated into a masked language modeling problem. Formally, let  $x = (x_0, x_1, \dots, x_n)$  be a raw paragraph with a length of  $n$  from the PCL dataset,  $T$  be a prompt template and  $x^T$  be the reformulated paragraph of  $x$  using  $T$ . For an example of SubTask 1 shown in Fig 1, with  $x = \text{“People ordered pizzas to be delivered, with the ample leftovers donated to local homeless shelters.”}$  and  $T = \text{“Is it patronizing or condescending? [MASK]”}$ , we wrap them into  $x^T$  as  $\text{“People ordered pizzas to be delivered, with the ample leftovers donated to local homeless shelters. Is it patronizing or condescending? [MASK]”}$ .

After the PCL problem is reformulated, an LM model  $M$  can be used to compute the probability of each word  $v$  in a vocabulary list to appear in the position of [MASK], denoted as  $P([\text{MASK}] = v|x^T)$ . To answer the cloze question  $\text{“Is it patronizing or condescending? [MASK]”}$ , if  $M$  has a large probability to fill [MASK] with some words of the meaning related to  $\text{“YES”}$ , then  $x$  is predicted as PCL class; if words having similar semantics to  $\text{“NO”}$  are more likely to be filled, then  $x$  is deemed as non-PCL class. In order to map the probabilities of predicted words to the probabilities of the labels  $\text{“YES”}$  and  $\text{“NO”}$ , we define a verbalizer to map a word  $v$ , from the label word set  $V_y$ , to a label  $y$ , form the label set  $Y$ . To construct the verbalizer,

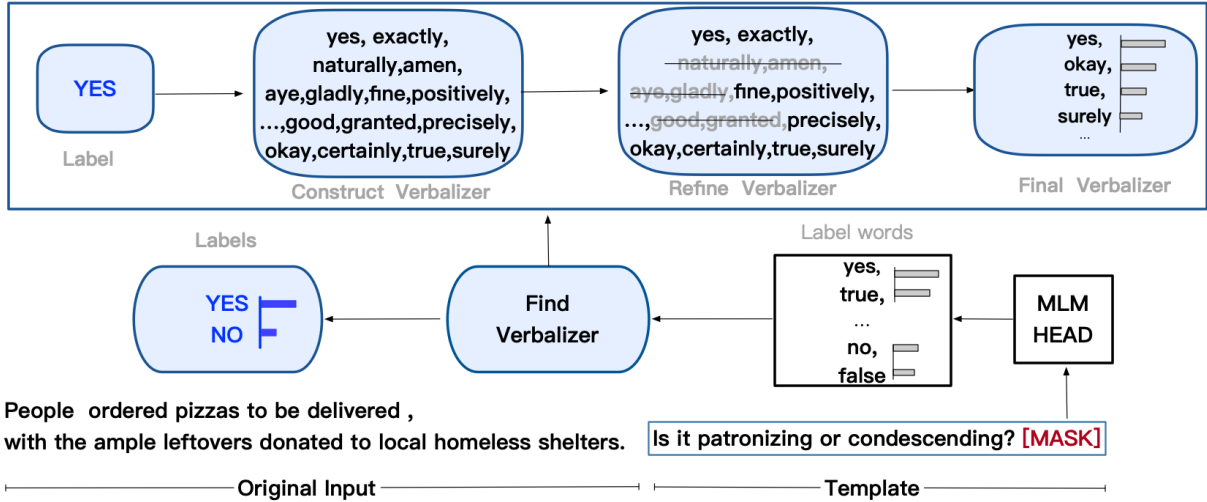


Figure 1: An example to illustrate our approach.

we first use synonym dictionaries to find the synonyms of labels as the candidate label words, and then we select the top-k candidate words, according to their usage frequency in public corpora, as the final label words. Assuming that all label words of the same label contribute equally to label prediction and the prior distribution of them is uniform, we use the mean of their predicted probabilities as the output probability of their corresponding label. Then, given a reformulated paragraph  $x^T$ , its predicted label  $\hat{y}$  can be obtained by

$$\hat{y} = \arg \max_{y \in Y} \left( \frac{1}{|V_y|} \sum_{v \in V_y} P([\text{MASK}] = v | x^T) \right)$$

With the above definitions, in the training process, a cross-entropy loss is applied to fine-tune the model  $M$  as our output model.

In SemEval-2022 Task 4, for the two subtasks, we use the same method to model and solve them. The only difference is the template used for each subtask. For the binary-classification subtask, only one template is provided to identify whether or not a paragraph is PCL class. For the multi-label-classification subtask, as there exist 7 PCL categories, each of them is given with a unique template for type detection.

### 3.2 Ensemble Strategy

To improve the robustness of the proposed approach, we take advantage of the cross-validation to make our model more stable and reliable. We randomly split the original PCL dataset into 10 parts, each of which has 1/10 samples of the dataset. According to the hold-one-out strategy, 10 folds of

the data can be acquired, and each fold has 9/10 samples for training and 1/10 samples for validation. During the training process, the best-trained model for each fold is saved, and the average output probability of all models is taken as the final prediction score. According to the score, each instance is assigned with a predicted label, and then the predicted labels are compared with the golden labels to compute the F1 score for validation. The best models acquired in the validation process are kept for the final online evaluation.

## 4 Experiment

### 4.1 Settings

The PCL dataset (Pérez-Almendros et al., 2020) containing 10,469 paragraphs is shared by both SubTask 1 and SubTask 2. The goal of SubTask 1 is to detect 993 PCL samples from the whole dataset, and the goal of SubTask 2 is to classify the 993 PCL samples into 7 specific PCL types. In our experiments, all trials are performed with Nvidia Tesla A100 and large-DeBERTa is used as the based pre-trained language model. Considering the small size of the task data, we take AdamW as the optimizer in the experiments with a learning rate of 1e-5 and a maximum epoch number of 10. Early-Stop strategy is also used in our training process. For other parameters, we set the batch size as 16 and the maximum sequence length as 256.

### 4.2 Used Strategies

For the sake of result analysis, we list all strategies used in our approach towards SemEval-2022 Task

Strategy	SubTask 1			SubTask 2		
	R	P	F1	macro-R	macro-P	macro-F1
CLS	61.7%	61.0%	61.4%	41.7%	41.1%	41.3%
Prompt	62.2%	61.8%	61.9%	44.3%	44.7%	44.4%
Prompt + Ensemble	63.0%	61.6%	62.3%	46.6%	46.3%	46.4%
Prompt + Ensemble + R-Drop	63.1%	62.0%	62.5%	46.7%	46.4%	46.5%
Prompt + Ensemble + R-Drop + EDA	63.3%	63.0%	63.1%	46.6%	48.2%	47.5%

Table 1: Experimental Results on SubTask 1 and SubTask 2

4 as the following.

- **Prompt:** Prompt strategy is the method described in Section 3.1.
- **Ensemble:** Ensemble strategy is the method described in Section 3.2.
- **CLS:** We directly feed the original text into DeBERTa model and apply a softmax layer on the CLS token for classification. This method is used as a comparison strategy, which does not reformulate the classification problem.
- **EDA:** (Wei and Zou, 2019) is a method of data augmentation, which introduces 4 operations, synonym replacement, random insertion, random swap, and random deletion, to boost performance on paragraph classification.
- **R-Drop:** (Wu et al., 2021) is a method to regularize dropout, which minimizes the bidirectional KL-divergence between the distributions of two sub-models sampled by dropout, to reduce model randomness.

### 4.3 Results and Analysis

Table 1 shows all strategy results with Precision, Recall, and F1. From the results, the effect of each strategy we used can be clearly observed. The last method including all Prompt, Ensemble, EDA and R-Drop performs best compared to other settings.

To show the superiority of the Prompt paradigm, we first compare it with CLS, which uses the traditional paradigm of classification. From the table, it can be seen that Prompt leads CLS by about 0.5% F1 score in SubTask 1 and 3.1% macro-F1 score in SubTask 2. As Prompt has natural advantages on few-shot learning and the size of the PCL dataset is small, such result is not out of expectation. With the limited training data, to avoid the occurrence of overfitting and improve the robustness of our

model, we then apply Ensemble strategy to the learning process. As shown in the table, with the robustness improved, a 0.4% increase of F1 score is achieved in SubTask 1. And for SubTask 2, the improvement is larger with a 2% increase of macro-F1 score, since the data size of each category in SubTask 2 is far less than 1/10 of that in SubTask 1. To further enhance model robustness, R-Drop is also adopted in our approach, which aims to reduce the randomness generated by the dropout module of network. Though the performance improvement of R-Drop is not big, the stability and convergence of our approach increases, making the results of repeated trials more similar.

Besides the techniques of modeling, we also try to improve the learning by increasing the quantity of the training data. As described in the previous sections, the size of the training data is small and PCL language is often written in an implicit style. Therefore, EDA strategy is performed to augment the PCL data, which replaces implicit words with their synonyms, and transforms the structure of paragraphs variously, to produce more data for training. With the augmented data, our approach achieves a 0.6% lift of F1 score on SubTask 1 and a 1% lift of macro-F1 score on SubTask 2.

## 5 Conclusion

In this paper, we propose a prompt-based learning approach for PCL detection, based on pre-trained language models. To reformulate the PCL detection problem into a masked language modeling problem, how our prompt template is designed have been fully discussed. We also introduce several other techniques, which have a positive impact on prompt-based learning. Through experimental comparison, our approach is proven as an effective solution to detect PCL language. As a result, in SemEval-2022 Task 4, our approach ranks 4th on the leaderboard of SubTask 1, and ranks 1st on the leaderboard of SubTask 2.

## References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.
- Danilo Croce, Giuseppe Castellucci, and Roberto Basili. 2020. Gan-bert: Generative adversarial learning for robust text classification with a bunch of labeled examples. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 2114–2119.
- 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*.
- Zhengjie Gao, Ao Feng, Xinyu Song, and Xi Wu. 2019. Target-dependent sentiment classification with bert. *Ieee Access*, 7:154290–154299.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Juanzi Li, and Maosong Sun. 2021. Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification. *arXiv preprint arXiv:2108.02035*.
- Thomas Huckin. 2002. Critical discourse analysis and the discourse of condescension. *Discourse studies in composition*, 155:176.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8018–8025.
- Thorsten Joachims. 1998. Text categorization with support vector machines: Learning with many relevant features. In *European conference on machine learning*, pages 137–142. Springer.
- Pradeep Kumar. 2020. Predictive analytics for spam email classification using machine learning techniques. *International Journal of Computer Applications in Technology*, 64(3):282–296.
- Branka Drljača Margić. 2017. Communication courtesy or condescension? linguistic accommodation of native to non-native speakers of english. *Journal of English as a lingua franca*, 6(1):29–55.
- Julia Mendelsohn, Yulia Tsvetkov, and Dan Jurafsky. 2020. A framework for the computational linguistic analysis of dehumanization. *Frontiers in artificial intelligence*, 3:55.
- Sik Hung Ng. 2007. Language-based discrimination: Blatant and subtle forms. *Journal of Language and Social Psychology*, 26(2):106–122.
- Carla Pérez-Almendros, Luis Espinosa-Anke, and Steven Schockaert. 2020. Don’t Patronize Me! An Annotated Dataset with Patronizing and Condescending Language towards Vulnerable Communities. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5891–5902.
- Carla Pérez-Almendros, Luis Espinosa-Anke, and Steven Schockaert. 2022. SemEval-2022 Task 4: Patronizing and Condescending Language Detection. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*. Association for Computational Linguistics.
- Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A Smith, and Yejin Choi. 2019. Social bias frames: Reasoning about social and power implications of language. *arXiv preprint arXiv:1911.03891*.
- Timo Schick and Hinrich Schütze. 2020a. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.
- Timo Schick and Hinrich Schütze. 2020b. It’s not just size that matters: Small language models are also few-shot learners. *arXiv preprint arXiv:2009.07118*.
- Jason Wei and Kai Zou. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196*.
- Lijun Wu, Juntao Li, Yue Wang, Qi Meng, Tao Qin, Wei Chen, Min Zhang, Tie-Yan Liu, et al. 2021. R-drop: regularized dropout for neural networks. *Advances in Neural Information Processing Systems*, 34.