

# Nowruz at SemEval-2022 Task 7: Tackling Cloze Tests with Transformers and Ordinal Regression

**Mohammadmahdi Nouriborji**

NODET, Iran\*

mohammadmahdinoori70@gmail.com

**Omid Rohanian**

Department of Engineering Science  
University of Oxford  
Oxford, UK

omid.rohanian@eng.ox.ac.uk

**David Clifton**

Department of Engineering Science  
University of Oxford  
Oxford, UK

david.clifton@eng.ox.ac.uk

## Abstract

This paper outlines the system using which team Nowruz participated in SemEval 2022 Task 7 “Identifying Plausible Clarifications of Implicit and Underspecified Phrases” for both subtasks A and B (Roth et al., 2022). Using a pre-trained transformer as a backbone, the model targeted the task of multi-task classification and ranking in the context of finding the best fillers for a cloze task related to instructional texts on the website Wikihow.

The system employed a combination of two ordinal regression components to tackle this task in a multi-task learning scenario. According to the official leaderboard of the shared task, this system was ranked 4th in both classification and ranking subtasks out of 21 participating teams. With additional experiments, the models have since been further optimised. The code used in the experiments is going to be freely available at <https://github.com/mohammadmahdinoori/Nowruz-at-SemEval-2022-Task-7>.

## 1 Introduction

Oxford dictionary defines cloze test as “a test of readability or comprehension in which a person is required to supply words which have been deliberately omitted from a passage” (Oxford University Press, 2022). In the context of NLP, a cloze format task is one in which the context is one or more sentences with masked spans and the model is expected to predict a suitable filler for each span. Cloze-format datasets have become popular in NLP recently as they are relatively easy to create automatically and provide high quality resources for model training (Rogers et al., 2021).

<sup>0</sup>National Organization For Development of Exceptional Talents

SemEval 2022 task 7 is framed as a cloze task in which the goal is to rank or classify fillers within a given context based on their suitability. The texts are taken from actual articles on an instructional website and the masked spans are placed at locations of edits made by users. We participated in this shared task in both ranking and classification parts and developed a transformer-based model that utilises both classification and regression components at the top layer. The code and the data used in these experiments are publicly available.

## 2 Related Work

Cloze tasks have been a subject of interest in Natural Language Understanding (NLU) in recent years, especially within the context of reading comprehension, story understanding, and summarisation (Deutsch and Roth, 2019; Sharma et al., 2018; Mostafazadeh et al., 2016). There is evidence that cloze tasks can be used to effectively pretrain or finetune language models in order to perform few shot learning (Schick and Schütze, 2021; Liu et al., 2021).

WikiHow is a community-edited open domain repository that hosts how-to articles on a variety of different subjects. It is possible to track edits made by users and compare different versions. There have been some recent computational works exploring this resource, including modelling of revision requirements (Bhat et al., 2020), and the effect of edits on fluency (Antonio and Roth, 2020) and vagueness (Debnath and Roth, 2021).

### 2.1 Masked Language Modelling

Masked Language Modeling (MLM) is a pre-training task which is widely used in transformer-based models. This task forms a self-supervised

cloze test by randomly removing some of the tokens of the sentence which will be further filled by the model during training. Prominent transformers including BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), DeBERTa (He et al., 2020) and T5 (Raffel et al., 2019) are trained using MLM as an auxiliary objective. Given the successes of transformers in most areas of machine learning and NLP, it is standard practice to fine-tune them for various down-stream tasks.

## 2.2 Ordinal Regression

Ordinal Regression (also known as Ordinal Classification) is a type of classification in which labels have order with respect to each other. Examples of ordinal regression tasks include age estimation (Niu et al., 2016), assessment of damage (Ci et al., 2019), and monocular depth estimation (Fu et al., 2018). In ordinal regression, performance of the model is sensitive to the order of the predictions with regards to the labels. For instance, in the age estimation task, the error of the model should be higher when it incorrectly predicts the age of 30 as 10, as opposed to when it predicts the age of 30 as 20. Ordinal Regression is commonly done by breaking the multi-class classification task into several binary classification subtasks within a multi-task learning scenario. The output of these binary classification subtasks should be rank-consistent to achieve good performance.

## 3 System Description

Figure 1 shows the overall architecture of our model. A pre-trained transformer sits at the base of the architecture. Different variations of this transformer are tried and reported in Sec 5. These include BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), DeBERTa (v1 & 3) (He et al., 2020, 2021), and T5<sup>1</sup> (Raffel et al., 2019) in different configurations.

At the bottom of the network, there are two components, one for regression and the other for classification. At training time, both of these are trained in tandem using ordinal regression and a combined loss. The loss function ensures that the ranks of the labels are kept consistent across these task (Sec. 3.6).

<sup>1</sup>Since T5 is originally an encoder-decoder language model, we only use its pre-trained encoder for our experiments

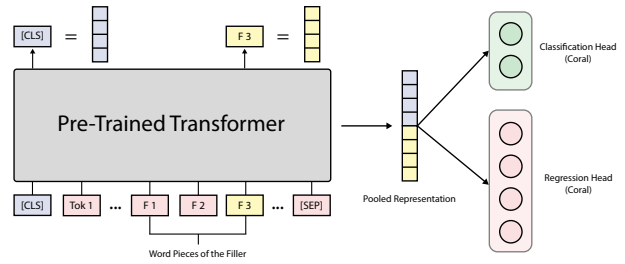


Figure 1: Overall architecture of the model. Note how a pooled representation is produced using the last token of the filler and the CLS token

### 3.1 Representing the Filler

To classify each filler, it is first placed in the location of the blank in the text to form a full context. The pre-trained transformer is equipped with a word-piece tokeniser that breaks down the filler to subword units and also uses a [CLS] token to represent the entire context<sup>2</sup>. To build a pooled representation for the filler, the contextualised representations for the last word piece and the entire context (i.e. representation of [CLS]) are concatenated and passed on to the next layer.

### 3.2 Multi-task learning with Ordinal Regression

Once a combined contextualised representation is obtained from concatenation of [CLS] and the filler, we address both subtasks of the shared task in a single multi-task learning architecture. The objective of this model is to predict both the class and the suitability score for each provided filler. The concatenated representation is fed to a fully connected feed-forward layer followed by a GELU activation function (Hendrycks and Gimpel, 2016). This layer projects the representation to a lower dimensional space. Subsequently, the output is passed on to two separate classification and regression heads.

### 3.3 Decomposing the Problem into Binary Classification tasks

As mentioned in Sec. 3, our model uses a special loss function named coral to perform ordinal regression. This loss function and the training procedure it requires are explained in Cao et al. (2020). Coral layers break a K-class classification problem into  $K - 1$  binary classification tasks as part of a multi-task learning scenario. For a coral layer with  $K - 1$

<sup>2</sup>T5 does not use CLS for context representation. In that case we just classify the last word-piece with no additional concatenation

units<sup>3</sup>, the value of loss is constructed from the sum of  $K - 1$  separate binary cross entropy losses belonging to each unit.

To train a coral layer, it is necessary to first transform the original labels to sets of binary labels. This is the step where the notion of order is introduced into the model. Given  $K$  classes, we convert each label to a collection of binary labels as follows:

$$f(y, k) = \begin{cases} 1 & \text{if } k < y \\ 0 & \text{if } k \geq y \end{cases} \quad (1)$$

$$Y_{ordinal}^{(i)} = \{f(y^{(i)}, k) | k \in \mathbb{W} \wedge k < K - 1\}$$

where  $y^{(i)}$  is the original label for the  $i$ th training instance and  $0 \leq y^{(i)} \leq K - 1$ .

### 3.4 Classification and Regression Heads

At the top end of the architecture there are two components for classification and regression. The classification head is a coral layer with two units which is used to address the classification subtask where we have three labels, namely, `Implausible`, `Neutral`, and `Plausible`. Note that the labels in the classification task have inherent order and for the model to be trained effectively, it is important that the training objective penalises misclassifications based on this underlying assumption. For instance, the error of the model should be more when it predicts an `Implausible` sample as `Plausible` compared to when it predicts an `Implausible` sample as `Neutral`. This is the motivation to use ordinal regression losses like coral.

The other head is assigned to the regression subtask and rates each filler for suitability. We converted continuous scores in the  $1 - 5$  range to labels with discrete values in the  $0 - 4$  range by either rounding or flooring the scores. For example, the original scores of  $\{1.333, 1.75, 2.5, 3.75, 4.25\}$  are mapped to  $\{0, 1, 2, 3, 3\}$  by rounding and  $\{0, 0, 1, 2, 3\}$  by flooring. We can frame the regression task as ordinal classification with five labels by binning the values. A 4-unit coral layer is used to perform this classification. These heads are jointly trained using a combined loss.

<sup>3</sup>A T-unit coral layer, is comprised of T binary classification units which share the same weights but have different biases.

### 3.5 Constructing Labels and Ranks from Heads

Since coral layers differ from normal dense layers, their output can not be directly converted to labels. Furthermore, the regression subtask is also framed as an additional classification task in our methodology. However, since the purpose is to report continuous scores rather than discrete labels, a unique conversion is necessary for the output of the regression head.

For the classification head, the goal is to output discrete labels in range of  $0 - 2$ . Given the output of the classification head for one sample  $\hat{C}$ , the conversion to the label is defined as follows:

$$f(y) = \begin{cases} 1 & \text{if } y > 0.5 \\ 0 & \text{if } y \leq 0.5 \end{cases} \quad (2)$$

$$c = \sum_{k=1}^2 f(\sigma(\hat{C}_k))$$

where  $c$  is the final label and  $\sigma$  is the sigmoid function.

For the regression head, the goal is to output continuous scores in range of  $1 - 5$ . Given the output of the regression head for one sample  $\hat{R}$ , the conversion to the continuous scores in range of  $1 - 5$  is defined as follows:

$$r = \left( \sum_{k=1}^4 \sigma(\hat{R}_k) \right) + 1 \quad (3)$$

### 3.6 Computation of Loss

Since we have three labels for the classification task, ordinal labels would be sets of binary values with the length of two. Given  $C^{(i)}$  as the converted classification label for the  $i$ th training sample and the  $\hat{C}^{(i)}$  as the output of the classification head for the  $i$ th training sample, we can define the classification loss as:

$$l_c^{(i)} = \sum_{k=1}^2 -C_k^{(i)} \log(\sigma(\hat{C}_k^{(i)})) - (1 - C_k^{(i)}) \log(1 - \sigma(\hat{C}_k^{(i)})) \quad (4)$$

For the regression task there are 5 different labels, and accordingly, ordinal labels would be sets of binary values with the length of four. Given  $R^{(i)}$  as the converted regression label for the  $i$ th training

sample and the  $\hat{R}^{(i)}$  as the output of the regression head for the  $i_{th}$  training sample we can define the regression loss in a similar way:

$$l_r^{(i)} = \sum_{k=1}^4 -R_k^{(i)} \log(\sigma(\hat{R}_k^{(i)})) - (1 - R_k^{(i)}) \log(1 - \sigma(\hat{R}_k^{(i)})) \quad (5)$$

Since we are aiming to perform both classification and regression, a joint loss is needed to combine losses from the two heads. For a given training batch of size  $n$ :

$$L_{total} = \frac{1}{n} \sum_{i=1}^n \lambda_c l_c^{(i)} + \lambda_r l_r^{(i)} \quad (6)$$

where  $\lambda_c$  is the weight associated with the classification loss and the  $\lambda_r$  is associated with regression.

## 4 Dataset and Experimental Setup

The shared task is based on [Anthonio et al. \(2020\)](#), where instructional texts from the website WikiHow<sup>4</sup> are used to create training instances for a cloze task. The deletions in the dataset are based on actual edits made by online users and they are assumed to represent certain types of clarifications to make a point more clear or disambiguate a sentence. Based on the actual edits, alternative fillers are automatically extracted and added to build a number of possible fillers. These texts along with the fillers were later annotated by humans and given plausibility scores from 1 to 5. There is also a separate type of annotation available in which there are 3 labels with discrete plausibility values. The shared task was organised in two separate tracks of classification and regression, depending on what kind of annotation was used for modelling the task. Table 2 & 3 show the basic statistics of the dataset in the shared task.

The dataset for this task consists of six features, named Resolved Pattern, Article title, Section header, Previous context, Sentence, and Follow-up context along with five different fillers for each sample. Resolved Pattern is one of the four following categories: IMPLICIT REFERENCE, ADDED COMPOUND, METONYMIC REFERENCE, and FUSED HEAD which indicates the relationship of the fillers with the context. Article title is the name of the article from which the paragraph is selected. Section

<sup>4</sup>A wiki-style online collection of how-to articles accessible at <https://www.wikihow.com>

header is the section from which the article is selected. Previous context is a few sentences before the sentence that contains the filler. Sentence is the sentence that contains the filler.

We obtained our best results when we used a custom formatting using which we can feed all the features to the model as textual input. Table 1 is an example of how we represent each training instance. Note how the ‘Text’ feature is constructed by concatenating previous and follow-up contexts with the target sentence.

### Example input while using all features

Resolved pattern: ADDED COMPOUND  
 Section header: Following a Basic Routine  
 Article title: How to Get Rid of Peeling Skin  
 Text: (...) 6. Never tear away loose skin. (...) 7. Protect your skin from sunlight. Exposure to direct sunlight can weaken your skin further and complicate the [Filler] problem. This is true regardless of whether your skin is peeling due to a sunburn or due to dryness.

Table 1: An example of how each input is formatted for training and inference. Each identifier followed by ":" represents a feature of the dataset represented in textual format.

Variation	implausible	neutral	plausible	total
train	5474 (27.40%)	7162 (35.85%)	7339 (36.75%)	19975 (100%)
dev	982 (39.28%)	602 (24.08%)	916 (36.64%)	2500 (100%)

Table 2: Statistics of the data for subtask A

Variation	1	2	3	4	5	total
train F	2254	4123	6259	5321	2018	19975
train R	1053	4421	4034	8441	2026	19975
dev F	645	458	481	596	320	2500
dev R	386	596	376	639	503	2500

Table 3: Statistics of the data for subtask B. R and F represent rounded and floored scores, respectively

## 5 Results and Discussion

The official test results for both subtasks are presented in Table 5. These are the best results that we have obtained on the test set prior to the end of the competition. On the official leaderboard of the shared task, our system was ranked 4th in both the ranking and classification tasks.

We have since performed extensive analyses on the effect of hyperparameters on different varia-

tions of our models, and since we do not have access to the true labels in the test set, we report our best results on the dev set. As can be seen in Table 4, we have produced our best results using DeBERTa-V3<sub>large</sub>. We have noticed two important factors that influence the final performance of the models. The first factor is batch size. Our best results were obtained on lower batch sizes of 8 and 16. The choice of rounded or floored numbers for subtask B is also a significant factor. The reason for this is that the distribution of labels changes depending on the normalisation method used.

## 6 Conclusion

In this work we developed a set of transformer-based models powered with ordinal regression to tackle an NLP cloze task as part of the SemEval 2022 shared task 7. The goal was to assign suitability scores or labels to several different provided fillers given each context and masked span.

Using a combined architecture based on ordinal regression that used training labels from both subtasks, we developed and trained models with a multi-task learning objective. The proposed system was ranked 4th out of 21 teams in both tracks of the shared task. In the subsequent analyses in the post evaluation phase, we have showed the effectiveness of this architecture in addressing this task. We compared different variations of our models and explored the effects of hyperparameters on model performance. The code and analyses are going to be publicly available.

## References

- Talita Anthonio, Irshad Bhat, and Michael Roth. 2020. [wikiHowToImprove: A resource and analyses on edits in instructional texts](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5721–5729, Marseille, France. European Language Resources Association.
- Talita Anthonio and Michael Roth. 2020. [What can we learn from noun substitutions in revision histories?](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1359–1370, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Irshad Bhat, Talita Anthonio, and Michael Roth. 2020. [Towards modeling revision requirements in wikiHow instructions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8407–8414, Online. Association for Computational Linguistics.
- Wenzhi Cao, Vahid Mirjalili, and Sebastian Raschka. 2020. Rank consistent ordinal regression for neural networks with application to age estimation. *Pattern Recognition Letters*, 140:325–331.
- Tianyu Ci, Zhen Liu, and Ying Wang. 2019. Assessment of the degree of building damage caused by disaster using convolutional neural networks in combination with ordinal regression. *Remote Sensing*, 11(23):2858.
- Alok Debnath and Michael Roth. 2021. [A computational analysis of vagueness in revisions of instructional texts](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 30–35, Online. Association for Computational Linguistics.
- Daniel Deutsch and Dan Roth. 2019. [Summary cloze: A new task for content selection in topic-focused summarization](#). 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)*, pages 3720–3729, Hong Kong, China. Association for Computational Linguistics.
- 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.
- Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. 2018. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2002–2011.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. [Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#).
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. In *International Conference on Learning Representations*.
- Dan Hendrycks and Kevin Gimpel. 2016. [Bridging nonlinearities and stochastic regularizers with gaussian error linear units](#). *CoRR*, abs/1606.08415.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.

Backbone Transformer	Model	Batch Size	Scores	Accuracy (subtask A)	Spearman’s rank (subtask B)
BERT	Bert <sub>base</sub>	8	R	57.80%	0.6020
	Bert <sub>base</sub>	8	F	57.48%	0.6038
	Bert <sub>base</sub>	16	R	56.68%	0.5950
	Bert <sub>base</sub>	16	F	57.92%	0.5960
	Bert <sub>large</sub>	8	R	57.60%	0.6323
	Bert <sub>large</sub>	8	F	59.12%	0.6341
RoBERTa	RoBERTa <sub>base</sub>	8	R	59.16%	0.6610
	RoBERTa <sub>base</sub>	8	F	58.68%	0.6582
	RoBERTa <sub>base</sub>	16	R	57.32%	0.6743
	RoBERTa <sub>base</sub>	16	F	59.44%	0.6695
	RoBERTa <sub>large</sub>	8	R	58.68%	0.6799
	RoBERTa <sub>large</sub>	8	F	60.20%	0.6928
DeBERTa-V3	DeBERTa-V3 <sub>base</sub>	8	R	61.72%	0.7194
	DeBERTa-V3 <sub>base</sub>	8	F	62.64%	0.7260
	DeBERTa-V3 <sub>base</sub>	16	R	63.12%	0.7287
	DeBERTa-V3 <sub>base</sub>	16	F	63.84%	0.7326
	DeBERTa-V3 <sub>large</sub>	8	R	<b>64.12%</b>	0.7406
	DeBERTa-V3 <sub>large</sub>	8	F	<b>64.12%</b>	<b>0.7411</b>
T5	T5-Encoder <sub>base</sub>	8	R	58.88%	0.6598
	T5-Encoder <sub>base</sub>	8	F	59.48%	0.6573
	T5-Encoder <sub>base</sub>	16	R	59.24%	0.6526
	T5-Encoder <sub>base</sub>	16	F	59.08%	0.6568
	T5 <sub>large</sub>	8	R	61.76%	0.6949
	T5 <sub>large</sub>	8	F	62.24%	0.6907

Table 4: Results on dev set for both subtasks, R and F represent rounded and floored scores respectively (as mentioned in Sec. 3.4)

Model	Batch Size	Scores	Accuracy (subtask A)	Spearman’s rank (subtask B)
RoBERTa <sub>large</sub>	8	F	61.00%	0.6700
DeBERTa-V1 <sub>large</sub>	8	F	61.00%	0.6900
T5-Encoder <sub>large</sub>	16	F	<b>62.40%</b>	<b>0.7070</b>

Table 5: Official results on leaderboard for both subtasks

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. *ArXiv*, abs/1907.11692.

Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. [A corpus and cloze evaluation for deeper understanding of commonsense stories](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849, San Diego, California. Association for Computational Linguistics.

Zhenxing Niu, Mo Zhou, Le Wang, Xinbo Gao, and Gang Hua. 2016. Ordinal regression with multiple output cnn for age estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Oxford University Press. 2022. [cloze](#), n.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits

of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.

Anna Rogers, Matt Gardner, and Isabelle Augenstein. 2021. Qa dataset explosion: A taxonomy of nlp resources for question answering and reading comprehension. *arXiv preprint arXiv:2107.12708*.

Michael Roth, Talita Anthonio, and Anna Sauer. 2022. SemEval-2022 Task 7: Identifying plausible clarifications of implicit and underspecified phrases in instructional texts. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*. Association for Computational Linguistics.

Timo Schick and Hinrich Schütze. 2021. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.

Rishi Sharma, James Allen, Omid Bakhshandeh, and Nasrin Mostafazadeh. 2018. [Tackling the story ending biases in the story cloze test](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 752–757, Melbourne, Australia. Association for Computational Linguistics.

## A Appendix

In this section, we are going to share the details of the hyperparameters used for the fine-tuning of our models and the final training procedure used for the submission.

### A.1 Hyperparameters

We fine-tuned all models for 5 epochs while keeping the embedding layers of all models frozen. We used the AdamW as our optimizer with a cosine learning rate scheduler from the Hugging Face library and a weight decay of 0.00123974 and an initial learning rate of  $1.90323e - 05$ . Also, the  $\lambda_l$  and  $\lambda_r$  (as mentioned in 3.6) are set to 0.5 in all of the reported experiments as shown in Table 6.

### A.2 Training Procedure for Submission

Once we found the best hyperparameters and models using Dev dataset, We used a combination of Training and Dev data to train our final models for submission. With this approach, we have been able to achieve a 3% improvement on accuracy and up to 0.05 improvement on Spearman’s rank correlation. Additionally, we found that combining Train and Dev data is noticeably less effective in the smaller models such as Bert<sub>base</sub> or RoBERTa<sub>base</sub>, however, this is a better strategy when it comes

Hyperparameters	
Epochs	5
Optimizer	AdamW
Learning Rate Scheduler	Cosine
Initial Learning Rate	1.90323e-05
Weight Decay	0.00123974
$\lambda_l$	0.5
$\lambda_r$	0.5

Table 6: Details of the Hyperparameters used for fine-tuning

to larger models such as RoBERTa<sub>large</sub> or T5<sub>large</sub>. Based on this, we suppose that with further tuning, DeBERTa-V3<sub>large</sub> can potentially surpass the state-of-the-art if it has access to the combination of Train and Dev data.