

Solomon at SemEval-2020 Task 11: Ensemble Architecture for Fine-Tuned Propaganda Detection in News Articles

Mayank Raj*
Ankita Gupta

Ajay Jaiswal*
Sudeep Kumar Sahoo

Rohit R.R
Vertika Srivastava

Yeon Hyang Kim

Samsung R&D Institute India, Bangalore

{ mayank.raj, ajay.jaiswal, rohit.r.r,
gupta.ankita, sudeep.sahoo, v.srivastava, purine.kim}@samsung.com

Abstract

This paper describes our system (Solomon) details and results of participation in the SemEval 2020 Task 11 "Detection of Propaganda Techniques in News Articles"(Da San Martino et al., 2020). We participated in Task "Technique Classification" (TC) which is a multi-class classification task. To address the TC task, we used RoBERTa based transformer architecture for fine-tuning on the propaganda dataset. The predictions of RoBERTa were further fine-tuned by class-dependent-minority-class classifiers. A special classifier, which employs dynamically adapted Least Common Sub-sequence algorithm, is used to adapt to the intricacies of repetition class. Compared to the other participating systems, our submission is ranked 4th on the leaderboard.

1 Introduction

In today's digital world, information is being shared more rapidly across boundaries owing to the rampant use of social media platforms. While these platforms provide a forum for people to exchange their views with ease, they also serve as a medium for spreading misinformation in the absence of proper quality control mechanisms. Two paradigms of misinformation are often reported, fake news and propaganda. While fake news corresponds to factually incorrect information, propaganda is more nuanced. It aims at advancing a specific agenda using psychological and rhetorical techniques (e.g use of biased or loaded language, repetition etc) and may not be always factually wrong.

Prior efforts have been made to segregate propagandist content (Rashkin et al., 2017), (Barrón-Cedeno et al., 2019), (Habernal et al., 2017) from non-propagandist. These methods use distant supervision based techniques to obtain document level labels, often labelling all articles from a propagandist news outlet as propaganda. However, such techniques introduce inevitable noise (Horne et al., 2018). To this end, (Da San Martino et al., 2019) have introduced a propaganda dataset where they label propaganda techniques at fragment level, thus providing finer control and explainability. For example, following is an instance of propaganda with type "Slogan".

Trump tweeted, "make America great again !"
Slogan

In this work we address the "Technique Classification" (TC) task whose objective is to identify the propaganda technique employed in a fragment of text surrounded by its context. In the above example, "make America great again" is the fragment in which *slogan* technique has been employed. It is surrounded by its context "Trump tweeted". The problem is modelled as a 14-class classification task. Our contributions in this work are summarized as follows :

- We employ transfer learning by fine-tuning Transformer based language model, RoBERTa (Liu et al., 2019) on the propaganda dataset.

*Equal Contribution

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

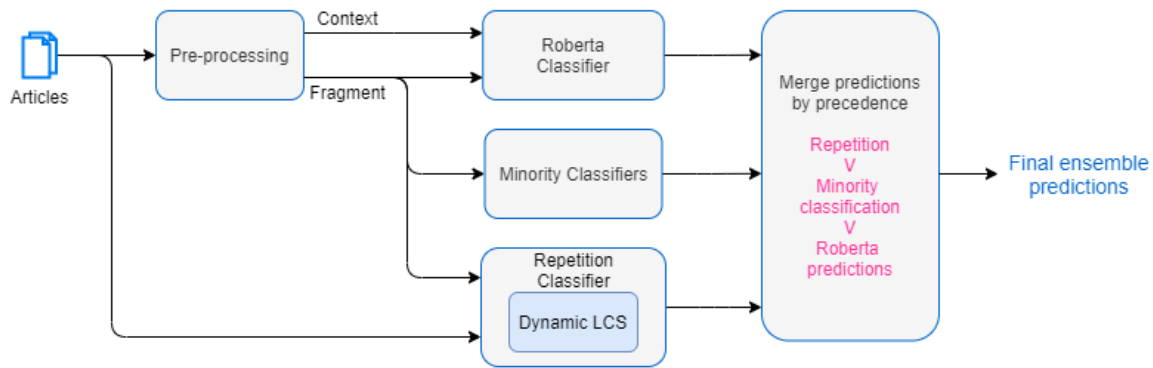


Figure 1: System Architecture with Precedence Order of Components

- We address the issue of minority class classification by designing ensemble of *one-versus-one (OVO)* classifiers, that vote for the presence/absence of the specific minority class.
- We handle the intricacies in the *Repetition* class by employing a novel algorithm based on dynamic least common sub-sequence approach.

The remainder of this paper is organized as follows: Section 2 describes the system details for identification of propaganda technique employed in a sentence. Section 3 describes our experiments and the evaluation results. Finally, we conclude in Section 4.

2 System Description

2.1 System Architecture

Our system comprises of three major components as depicted in Fig 1. The first component is fine-tuned RoBERTa that gives predictions for 14 classes. The predictions of RoBERTa are adjusted based on the output of Minority Classifiers followed by the Repetition classifier. If the Minority and Repetition classifier refrain from making a prediction, we retain the predictions of RoBERTa. The order of precedence among these three components is depicted in the architecture diagram.

2.2 Fine Tuned RoBERTa

Our first system component comprises of fine-tuning (end-to-end) a pre-trained language model on the downstream task of propaganda technique classification. We leverage RoBERTa, a transformer architecture based language model for fine-tuning. RoBERTa uses BERT’s (Devlin et al., 2018) masked language modelling strategy with modified hyper-parameters and is trained for longer period of time on substantially larger dataset (10 times more in size) in comparison to BERT.

In order to fine-tune RoBERTa for 14-class TC task, we modify the last layer, comprising of 14 hidden nodes with softmax activation. This layer is specifically trained only on the downstream TC task, whereas all other layers of RoBERTa are fine-tuned.

We experimented with three methods of input, i.e., *‘fragment only’*, *‘context only’* and *‘fragment and context both’*. While the fragments are mainly decisive in classification, they face the issue of shorter length and missing background information, hence the reason for sub-optimal performance of *‘fragment only’* approach. Using the *‘context only’* approach, led to overshadowing the emphasis on the fragment embedded in a long surrounding text. Thus, the combined approach of using both fragment and context worked best as illustrated in our experiments section.

We input both fragment and context as an input to RoBERTa by concatenating them together along with the separator to mark the beginning and end of each sequence. This is in similar lines to how RoBERTa takes input for a pair of sentence for tasks such as QA, NLI etc (Liu et al., 2019). Using fragment and context as a pair of input sentences to RoBERTa, helps us leverage bidirectional cross attention between

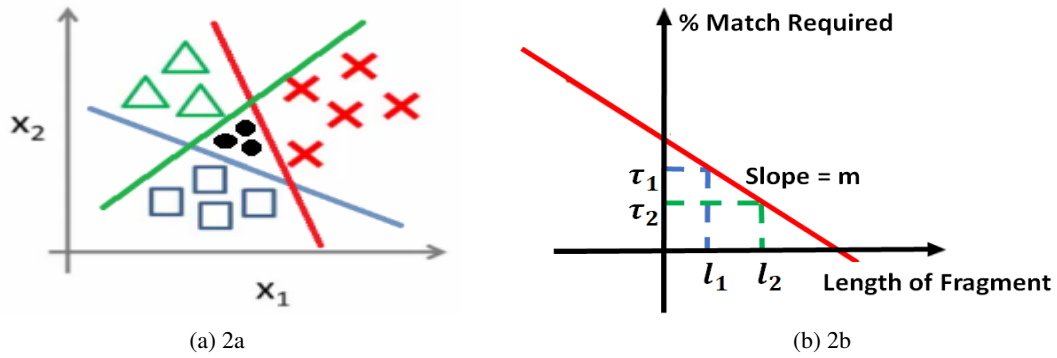


Figure 2: (a) : Ensemble of One-Versus-One Classifiers Segregating Minority Class. (b) : Dynamic-LCS depicting the variation between % match required and length of fragment.

the two sentences. The predictions of RoBERTa are then adjusted further based on the Minority and Repetition Classifiers as described next.

2.3 Minority Classifiers

The dataset has severe class imbalance owing to the varied frequency of occurrence of the different types of propaganda techniques in real life. For few minority classes, the number of positively labelled data samples are extremely small in comparison to the other classes. We tackle each one of these minority classes namely, “*Bandwagon, Reductio ad Hitlerum*”, “*Appeal to Authority*”, “*Black and White Fallacy*”, “*Whataboutism, StrawMen, RedHerring*”, and “*Thought Terminating Cliche*” by training 5 separate hierarchical classifiers.

We denote these 5 classifiers as *level-1* classifier in the hierarchy. Each one of these *level-1* classifier is further composed of an ensemble of $n-1$ (13 in our case) one-versus-one linear classifiers, denoted as *level-2* classifiers. We then take vote of all these $n-1$ classifiers and aggregate their predictions to obtain the final prediction confidence of a *level-1* classifier, denoting its confidence in predicting a corresponding minority class. If the confidence is above a threshold, we treat it as a positive example of that minority class. Rationale behind using such an approach is depicted in Fig 2a, where learning three OVO classifier can clearly help segregate the minority class (black) from others. Also, this ensemble based classifier essentially ensures that each one of the $n-1$ classifiers should, with very high confidence, predict the positive class. This helps us in classifying these minority classes within the data restricted settings and at the same time overcome the issue of possible over-fitting in such scenarios. We verify this intuitive explanation in our experiments.

A *level-2* classifier is a simple linear classifier and hence lends itself reduced time and computational complexity as an added benefit. The final predictions of the *level-1* minority class classifiers are used to over-rule the predictions of RoBERTa. In cases where the prediction confidence is less than the threshold, predictions of RoBERTa take precedence.

2.4 Repetition

Repetition is another class that we handle separately based on our analysis of the data. Repetition class, by definition, refers to repeating the same phrase over and over again across the message to enhance its impact on the audience. We formulate this problem as detecting the presence of dynamic least common sub-sequence (dynamic-LCS) between the fragment and the context. We avoid using exact match as fragments are repeated across the message with some minor modifications. For instance, consider the fragment “*ammunition was purchased under someone else’s name*”. This fragment is repeated in various forms across the message, such as “*ammo was bought under someone else’s name*”.

In order to measure the presence of LCS, and hence repetition, we detect if the % match between fragment and context is greater than a threshold (τ). This threshold is dynamically adapted based on the length of the fragment as shown in Fig 2. If the length of the fragment is small, high LCS is desired for a

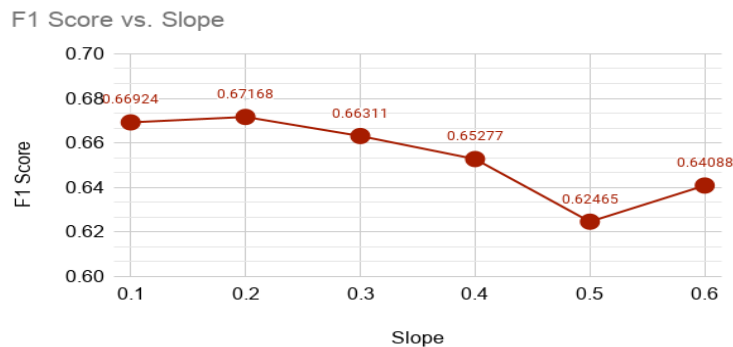


Figure 3: Variation of F1 score with respect to slope hyperparameter.

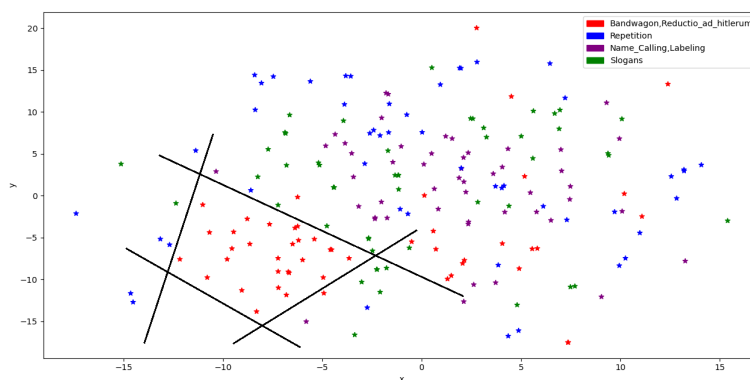


Figure 4: t-SNE plot of RoBERTa embedding

reliable indication of sufficient match and vice-versa. We model this by tuning the slope parameter (on dev-set) of the curve between length of the fragment and desired match threshold. Mathematically, τ is defined as below.

$$\tau = 100 - (m) * l \quad (1)$$

For a given slope of $m=0.2$, a string of length 100 will need 80% match. This implies that at-least 8 characters out of 10 should match for a string length of 10.

3 Experiments and Analysis

3.1 Analysis of fine-tuned RoBERTa

Table 1 mentions our experimental results obtained for various models. As a baseline, we also experimented with an SVM classifier (linear and RBF kernel) trained on pre-trained BERT embeddings of propagandist fragment. Also, for brevity we also checked the performance of fine-tuned BERT which was trained end-to-end similarly to how we trained RoBERTa. Our Fine-tuned RoBERTa gave the best performance with an an F1 score of 0.64, being pre-trained on remarkable larger dataset in comparison to fine-tuned BERT. Further, adapting the predictions of RoBERTa based on minority classifiers and Repetition classifier led to F1 scores of 0.65 and 0.67 respectively.

We used uncased base model of BERT with batch size of 16, maximum sequence length of 512, weight decay of 0.01 and trained using adam optimizer with a learning rate of $3e-05$ for 5 epochs. For RoBERTa, we used uncased large model with a batch size of 8, maximum sequence length of 512 and weight decay of 0.1. It was trained using Adam Optimizer with learning rate of $1e-05$ for 10 epochs.

Models	F1 Score
Official Baseline (Logistic Regression + Sentence length)	0.25196
Pre-trained BERT embedding + SVM (Linear Kernel)	0.44120
Pre-trained BERT embedding + SVM (RBF Kernel)	0.48636
Fine-tuned BERT	0.59548
Fine-tuned Roberta	0.63891
Roberta + Minority Classifier	0.65744
Roberta + Minority Classifier + Repetition Classifier	0.67168

Table 1: Experiment results of different model for the 14 class propaganda detection

Techniques	RoBERTa F1
Loaded Language	0.78655
Name Calling, Labeling	0.74176
Repetition	0.23015
Flag Waving	0.74359
Exaggeration, Minimisation	0.50746
Doubt	0.61224
Slogans	0.71605
Appeal to fear-prejudice	0.41509
Causal Oversimplification	0.31250
Appeal to Authority	0.02809
Black & White Fallacy	0.01683
Whataboutism, StrawMen, RedHerring	0.11432
Thought-terminating-Cliches	0.27838
Bandwagon, Reductio-Hitlerum	0.00000
Total F1	0.63891

Table 2: F1 Score of 14 classes with Fine-tuned RoBERTa Model.

3.2 Analysis of Minority Classifiers and Repetition Classifier

As depicted in Table 1, predictions of fine-tuned RoBERTa perform well on all the majority classes. However, it exhibits poor F1 score for minority classes. We believe that high class imbalance in the dataset resulted in such an observation and motivated us to look closely at data distribution across classes. To illustrate our intuition introduced in Section 2.4, we projected the sentence embeddings (RoBERTa) of minority samples onto a 2D subspace using t-Distributed Stochastic Neighbor Embedding (t-SNE) plot. Fig 4 depicts how the data points corresponding to the minority class “*Bandwagon, Reductio_ad_hitlerum*” are arranged relative to other majority class samples. It can be observed that the minority samples in this case can be segregated well by learning an ensemble of linear classifiers, thus reinforcing our approach.

In experiments, we use $n_components=2$, $perplexity=40$, $n_iter=300$ while plotting t-sne plots. We also employ oversampling with replacement while training OVO *level-2* classifiers. We kept the threshold for aggregated confidence of all *level-2* classifiers as 0.95, above which a *level-1* classifier will vote in favour of a minority class. As presented in Table 1 Minority classifiers led to a 2% boost in the F1 score. Specifically, the F1 score of *Bandwagon, Reductio_ad_hitlerum* increased from 0.0 to 0.89 on the dev set.

Further improvement in F1 score can be observed due to Repetition Classifier in Table 1. We also measure how the system performance varies with the change in hyper-parameter slope “ m ”. Fig 3 depicts that slope=0.2 gives best performance with highest F1-score of 0.67.

4 Conclusion

In this paper, we examine capability of the transformer based pre-trained language model, RoBERTa. We illustrated that Fine-tuned RoBERTa performed better on minority classes compared to BERT. We

also introduced novel approach of handling the minority classes using an ensemble of one-vs-one simple classifiers. Furthermore, we handle the Repetition class separately using Dynamic LCS algorithm. Experiments show the improvement in F1 score when fine-tuned RoBERTa predictions are augmented with minority class classifiers and repetition classifier.

References

- Alberto Barrón-Cedeno, Giovanni Da San Martino, Israa Jaradat, and Preslav Nakov. 2019. Proppy: A system to unmask propaganda in online news. volume 33, pages 9847–9848.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeno, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news article. 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 5640–5650.
- Giovanni Da San Martino, Alberto Barrón-Cedeño, Henning Wachsmuth, Rostislav Petrov, and Preslav Nakov. 2020. SemEval-2020 task 11: Detection of propaganda techniques in news articles. In *Proceedings of the 14th International Workshop on Semantic Evaluation, SemEval 2020, Barcelona, Spain, September*.
- 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*.
- Ivan Habernal, Raffael Hannemann, Christian Pollak, Christopher Klamm, Patrick Pauli, and Iryna Gurevych. 2017. Argotario: Computational argumentation meets serious games. *arXiv preprint arXiv:1707.06002*.
- Benjamin D Horne, Sara Khedr, and Sibel Adali. 2018. Sampling the news producers: A large news and feature data set for the study of the complex media landscape. In *Twelfth International AAAI Conference on Web and Social Media*.
- 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 preprint arXiv:1907.11692*.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2931–2937.