

# BPGC at SemEval-2020 Task 11: Propaganda Detection in News Articles with Multi-Granularity Knowledge Sharing and Linguistic Features based Ensemble Learning

Rajaswa Patil<sup>1</sup> and Somesh Singh<sup>2</sup> and Swati Agarwal<sup>2</sup>

<sup>1</sup>Department of Electrical & Electronics Engineering

<sup>2</sup>Department of Computer Science & Information Systems

BITS Pilani K. K. Birla Goa Campus, India

{f20170334, f20180175, swatia}@goa.bits-pilani.ac.in

## Abstract

Propaganda spreads the ideology and beliefs of like-minded people, brainwashing their audiences, and sometimes leading to violence. SemEval 2020 Task-11 aims to design automated systems for news propaganda detection. Task-11 consists of two sub-tasks, namely, Span Identification - given any news article, the system tags those specific fragments which contain at least one propaganda technique and Technique Classification - correctly classify a given propagandist statement amongst 14 propaganda techniques. For sub-task 1, we use contextual embeddings extracted from pre-trained transformer models to represent the text data at various granularities and propose a multi-granularity knowledge sharing approach. For sub-task 2, we use an ensemble of BERT and logistic regression classifiers with linguistic features. Our results reveal that the linguistic features are the reliable indicators for covering minority classes in a highly imbalanced dataset.

## 1 Introduction

Propaganda is biased information that deliberately propagates a particular ideology or political orientation (Aggarwal and Sadana, 2019). Propaganda aims to influence the public’s mentality and emotions, targeting their reciprocation due to their personal beliefs (Jowett and O’donnell, 2018). News propaganda is a sub-type of propaganda that manipulates lies, semi-truths, and rumors in the disguise of credible news (Bakir et al., 2019). The emphasis on this manipulation differentiates propaganda and its various classes from each other and free verbalization (Hackett, 1984). News propaganda can lead to the mass circulation of misleading information, shared agenda, conflicts, religious or ethnic reasons, and can further even lead to violence and terrorism. Due to massive size, high velocity, rich online user interaction, and diversity, the manual identification of propaganda techniques is overwhelmingly impractical. Hence, the development of a generalized system for propaganda detection in news articles is a vital task for security analysts and society (Barrón-Cedeno et al., 2019).

In *SemEval2020-Task11*, Da San Martino et al. (2020) propose a corpus of 550 news articles for propaganda detection. Each article is annotated with propaganda spans belonging to 14 propaganda techniques. The annotation is performed at the fragment level. The task of propaganda detection is divided into two sub-tasks; Span identification (sequence labeling) and technique classification (sequence classification). Span identification sub-task aims to detect the propagandist spans of text in the news articles. Whereas, the technique classification sub-task aims to classify the propaganda spans into various propaganda techniques. The work presented in this paper aims to provide independent approaches for both sub-tasks 1 and 2. Most recent approaches for propaganda detection task use pre-trained transformer models. This work propose a multi-granularity knowledge sharing model built on top of the embeddings extracted from these transformer models for the propaganda span detection sub-task. Further, we show the effectiveness of ensembling linguistic features-based machine learning classifiers with these transformer models covering the minority-classes for the technique classification sub-task.

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

## 2 Background

Before the ascension of the detection of fake news and propaganda in NLP, Mihalcea and Strapparava (2009) introduced the automatic detection of lying in text. In this work, the authors proposed three datasets based on *abortion*, *death penalty*, and *best friends* and approached the multi-class classification problem. The Amazon Turk Service performed the annotations for these datasets. They used support vector machine and naive bayes classifiers for conducting their experiments. On the other hand, Ciampaglia et al. (2015) had employed knowledge graphs for computational fact-checking and utilizing data amassed from Wikipedia. Detection of propaganda in news articles was advocated by Rashkin et al. (2017) and Barrón-Cedeno et al. (2019). The former proposed the use of Long Short-term Memory (LSTM) models and machine learning methods for deception detection and classification to different types of news, *trusted*, *satire*, *hoax*, and *propaganda*. The latter presented *proppy*<sup>1</sup> the first publicly available real-world and real-time propaganda detection system for online news. Fake news and propaganda have been getting more attention recently (Bourgonje et al., 2017; Helmstetter and Paulheim, 2018; Jain and Kasbe, 2018). Previous research has leveraged BERT (Devlin et al., 2018) to extract features for fake news detection (Singhal et al., 2019). The constrained resources and lack of annotated corpora are considered to be an immense challenge for researchers in this field. Da San Martino et al. (2019) proposes a new dataset with text articles annotated on the fragment level with one of 18 given propaganda techniques. To classify them, they also employ BERT models for this high granularity task. Given the relatively low count of some of these 18 techniques, they merged similar underrepresented techniques into superclasses like *Bandwagon*, *Reductio\_ad\_hitlerum*, *Exaggeration*, *Minimisation*, *Name\_Calling*, *labelling*, and *Whataboutism*, *Straw\_Men*, *Red\_Herring* and eliminate *Obfuscation*, *Intentional Vagueness*, *Confusion* to compile 14 classes (Da San Martino et al., 2020).

## 3 Methodology

In this section, we explain our approach and the road-map to our conclusion. A detailed explanation of our implementations is provided in Section 4. Inspired by Da San Martino et al. (2019), we consider three granularities of propaganda detection: Document-level, Sentence/Fragment-level, and Token-level. The Span Identification sub-task focuses on high-granularity propaganda detection through a token-level sequence labeling task. Whereas, the Technique Classification sub-task focuses on sentence/fragment level classification. Even though both the sub-tasks are restricted to a single granularity, the context from other granularities can be used by defining auxiliary objectives for each granularity.

### 3.1 Sub-Task 1: Span Identification

We approach the span identification sub-task with token-level binary classification. We label every token as *belongs / does not belong* to a propaganda span. A continuous sequence of propaganda tokens in the text combine to form a propaganda span. We use pre-trained word embeddings to represent the tokens. As the context from lower granularity is essential during the classification, using only token-level information for this task is not sufficient. To incorporate the lower granularity context for token-classification, we propose a multi-granularity knowledge-sharing model. We use two lower granularity contexts for this: the sentence and the article (document) to which the token belongs. Figure 1 illustrates the high-level architecture of the proposed model for span identification. We train the proposed model simultaneously for three different objectives (one for each granularity): article-level regression, sentence-level classification, and token-level classification.

#### 3.1.1 Article-level Regression

For each article in the dataset, we have the article-title and the article-content split at sentence-level. The title of an article is a good summarized representation of the article’s content. We take the mean-pooled embedding of all the word embeddings from the article’s title to get the title embedding ( $Emb_A$ ). We represent each article by its respective  $Emb_A$ . Further,  $Emb_A$  is passed through an article encoder (Figure 1) to get a lower-dimensional article encoding ( $Enc_A$ ). For every article, we calculate the

<sup>1</sup>Proppy, propaganda detection system. <https://proppy.qcri.org/>

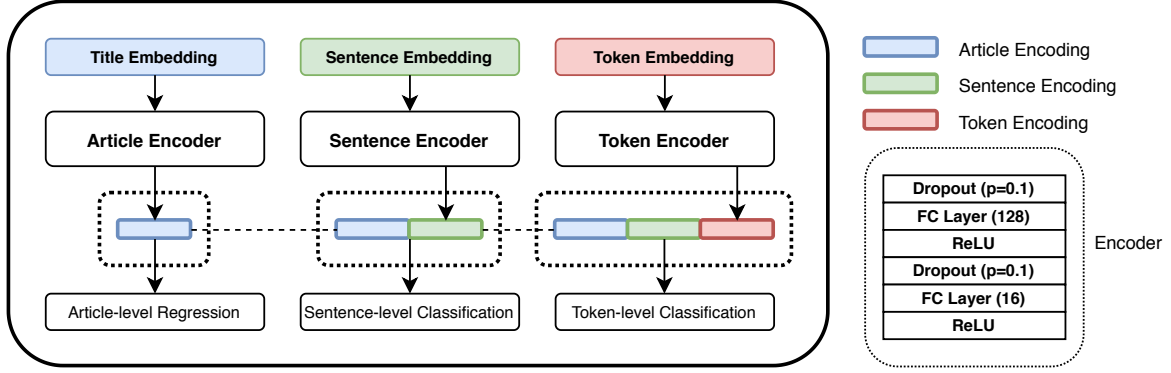


Figure 1: Model Architecture for the Span Identification sub-task.

normalized count of sentences with propaganda spans. We refer to this value as the propaganda-ratio ( $Ratio_P$ ) which is defined according to Equation 1.  $Ratio_P$  is restricted to the range  $[0, 1]$ .  $Enc_A$  is used for regression over the  $Ratio_P$  values of the articles. To minimize the error in predicted  $Ratio_P$  values, we use Smooth L1 Loss ( $SL1_A$ ) as an objective function for this regression task to avoid exploding gradients from the outlying predictions.

$$Ratio_P = \frac{\text{Number of sentences with propaganda spans}}{\text{Total number of sentences in the article}} \quad (1)$$

### 3.1.2 Sentence-level Classification

We label every sentence in the article as *contains / does not contain* a propaganda span. For every sentence, we obtain its sentence embedding ( $Emb_S$ ) by taking the mean-pooled embedding of all the word embeddings from the sentence.  $Emb_S$  is passed through a sentence encoder (Figure 1) to get a lower dimensional sentence encoding ( $Enc_S$ ). We concatenate  $Enc_S$  with its corresponding  $Enc_A$  to incorporate the article context. The concatenated encoding is further used for the sentence-level binary classification task. We use Binary Cross-Entropy Loss ( $BCE_S$ ) as an objective function for this task.

### 3.1.3 Token-level Classification

For token-level binary classification, we pass the token word embedding ( $Emb_T$ ) through a token encoder (Figure 1) to get a lower-dimensional token encoding ( $Enc_T$ ). We concatenate  $Enc_T$  with its corresponding  $Enc_S$  and  $Enc_A$ . This concatenated encoding is finally used for the token-level binary classification task. We use Binary Cross-Entropy Loss ( $BCE_T$ ) as an objective for this task. The concatenation of encoding with its corresponding lower granularity encoding incorporates a unidirectional knowledge transfer from *lower to higher granularity* level. To implement an implicit bi-directional knowledge transfer, we perform simultaneous training for tasks across all the granularities. To perform the training, we define a combined multi-granularity objective function ( $Objective_{MG}$ ) as shown in Equation 2. We establish a trade-off between the Smooth L1 Loss ( $SL1_A$ ) and the classification losses ( $BCE_S$  and  $BCE_T$ ) with  $\alpha$  as the trade-off factor. Model parameters across all three granularities are trained together to optimize  $Objective_{MG}$  objective function, resulting in a multi-task learning setting (Figure 1).

$$Objective_{MG} = (\alpha * SL1_A) + (1 - \alpha) * (BCE_S + BCE_T) \quad (2)$$

## 3.2 Sub-Task 2: Technique Classification

This sub-task is a multi-class classification of propaganda techniques. Our system classifies each propagandist statement amongst the 14 given techniques. The dataset is highly imbalanced, and there is a lot of variation in the length of sequences. Therefore, we ensemble our machine learning algorithm and deep learning transformer architecture to provide a solution with better generalization across the classes.

### 3.2.1 Machine Learning

For hand-crafted feature-based machine learning methods, we preprocess the sentences and extract their pragmatic and lexical features. Machine learning algorithms such as XGBoost and Logistic Regression use these features and perform better than the given baseline substantially and generalize well. Since news articles are written in a formal language with very few typographical or linguistic errors, identifying these features is less technically challenging compared to free-form text.

We perform preprocessing on our text to remove any unwanted content and enrich our text for feature extraction. The various procedures for the preprocessing include UTF-8 conversion, removal of non-ASCII characters, lower casing, lemmatizing/stemming, and removal of extra white-spaces, newlines, and stop-words. We create the feature space for the classification using three metadata types: contextual (sequence length, count of '!' and '?'), content (word and char n-gram TFIDF, part of speech) and context-based metadata (polarity of a sentence<sup>2</sup>, and topic modeling). However, these models do not effectively capture the sequential nature of the data.

### 3.2.2 Deep Learning : Stacked LSTMs, Pretrained Embeddings, CNN-LSTMs

We shift our focus to deep learning models that retain the sequential information by training Long-Short-Term-Memory [LSTM] Networks (Gers et al., 1999). The non-linear decision boundaries of these neural networks capture the complex features inside text. We convert the sentences to lowercase, non-ASCII characters are removed and tokenized to words. While validating or testing all obscure words are converted to a special [UNK] (Unknown) token. Approximately 64% of the words present in the development set's vocabulary intersect with the training set's vocabulary. Also, 85.8% of the words in the training-vocabulary have a count lesser than five.

Each word is encoded to fixed-size-vector or embedding. These embeddings are randomly initialized and passed to a stacked LSTM network with a classification layer on top. These networks give us comparable results. For this approach, we use the stacking of two LSTM layers with 256 hidden-sizes each. On increasing the number of stacked layers and the hidden-size of each LSTM unit, the models tend to overfit and neglect the minority classes due to the relatively small feature space and training samples per class. Therefore to improve the semantic relationship between the words and increase the training vocabulary, we propose systems that have prior knowledge of the language and a larger vocabulary set. Such systems show improvement on these linguistically rich corpora over models with randomly initialized embeddings (Erhan et al., 2010). We can infer this because their language modelling is more refined and wider.

To increase the vocabulary for our model and to extract the contextual importance of words, we incorporate pre-trained word embeddings such as "Global Vectors" [GloVe] (Pennington et al., 2014). These embeddings map each input token to a fixed-length vector. We set them to be non-trainable to avoid changing these vectors substantially. Pre-trained embeddings help the model to understand the underlying relationships.

Since GloVe is trained on a large corpus to obtain vector representation of words, 98.8% of the vocabulary in the development set and 97.9% of the training vocabulary intersects with its own. Therefore, in total less than 2% of these words are out of vocabulary and address the issue of vocabulary coverage in above models. On inspection, we find that most of the out-of-vocab words in the GloVe are conjunct words commonly used in the text, such as *jawboned*, *antichrists*, *trumpists*, *cyberspies*, *atleast*. Therefore we use sub-word tokenization and break each of these absent words into two in-vocab words where the smaller word is at least two characters long like *jaw boned*, *anti christ*s, *trump ist*s, *cyber spies*, *at least*. The sub-word tokenization increases our coverage to 99.1% with the remaining words either being proper-nouns or non-alphabetic patterns. Since these words convey a little meaning, and are removed during training and testing.

We employ CNNs (Convolutional Neural Networks) and CNN-LSTM models as well by adding convolution filters before extracting and pooling information from the embeddings. CNN-LSTM models help to extract a sequence of higher-level phrase representations (Zhou et al., 2015). As a result of reduced

---

<sup>2</sup>TextBlob API. <https://textblob.readthedocs.io/en/dev/>

inputs for the LSTM layers, these models converge faster (in less number of epochs) and quicker (in terms of training time per epoch), but the performance does not improve.

We pad all the sequences which have less than 100 tokens and truncate the others to the maximum length of 100. The tokens are encoded to word embeddings. The obtained embeddings are fed to the stacked LSTM layers, with dropout regularization to prevent over-fitting. We pass the LSTM layers' outputs to the linear classification layers with sigmoid activation to obtain the categorical cross-entropy loss. We use Adam optimizer with a small learning rate to update the parameters for 50 epochs (Kingma and Ba, 2014).

### 3.2.3 Deep Learning : Transformer Architectures

Consequently, we transit to deep pre-trained transformer architectures with attention mechanism (Vaswani et al., 2017). We employ BERT-base (Devlin et al., 2019) model which has 12 Layers. BERT tokenizer implements word-piece tokenization, which eliminates out-of-vocabulary words by splitting the words into sub-words. For example the word "judgmental" is broken to ["judgement", "##al"]. This model has been pre-trained on large corpus for masked language modelling and hence can be used for many downstream NLP tasks such as classification, question answering systems, and named-entity-recognition. We freeze most of the layers of the model and train the remaining (layer 10, 11, 12) with a very small learning-rate:  $1e-4$  for four epochs with a classification layer on top of the extracted pooled output embeddings.

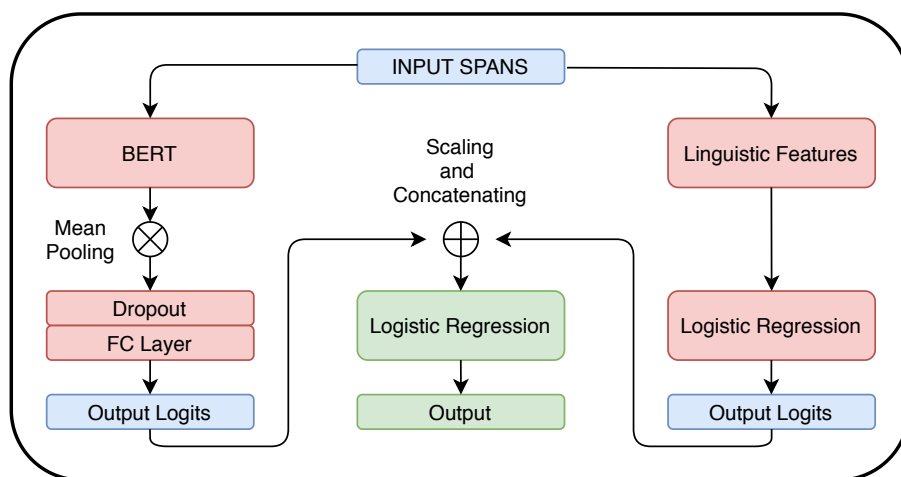


Figure 2: Ensemble Model Architecture for Technique Classification sub-task.

### 3.2.4 Ensemble Learning

The majority of the deep learning models perform better than the machine learning models on the validation split. On closer inspection, we observed that deep learning models tended to neglect the minority classes and had some unpromising class-wise F1 scores (discussed in Section 5.2). On the other hand, machine learning models had no zero-scores for any technique but proved to be less decisive for the majority-classes. We infer that training an ensemble of these models; the proposed architecture can improve in terms of generalization across all the classes. In this manner, we account for the gain in overall performance and also avoid zero-scores for the minority classes. Therefore we concatenate the scaled outputs from both the models and pass this to a logistic regression classifier. We penalize this classifier by the L2-norm (euclidean distance) and keep class wise weights for the objective in inverse proportion to their count in training set to obtain the final predictions. Our F1 scores for minority classes show massive improvements while only a minor change (+0.006) in the overall micro-F1 score.

## 4 Experimental Setup

### 4.1 Sub-Task 1: Span Identification

For the span-identification sub-task, we use word embeddings from pre-trained transformer models (Vaswani et al., 2017), extracted using *flairNLP*<sup>3</sup> (Akbik et al., 2018). We test our model with various embeddings such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), XLNet (Yang et al., 2019) and GPT-2 (Radford et al., 2019), which differ in their model architectures and pre-training methods. These embeddings and their respective transformer-models are kept non-trainable throughout our experiments so that the standalone performance of our proposed system can be monitored effectively. We split the given training dataset at token-level with a 95-5 split for training and validation purposes, respectively. We compute the accuracy, F1 score, precision, and recall metrics for the token-level binary-classification. Since there is a high class-imbalance for token-level classification, we consider the token-level F1 score for validation purposes. The official evaluation metrics (span-level F1, precision, and recall) used for this task are calculated with respect to the overlaps between the predicted and the ground-truth spans (Da San Martino et al., 2019).

### 4.2 Sub-Task 2: Technique Classification

For this sub-task, we use BERT (base-uncased variant), a pre-trained transformer model based on the work done by Devlin et al. (2018) with a classification layer on top. We extract the contextual pooled output embedding from the BERT model using *transformers*<sup>4</sup>, a python package by Hugging Face. All layers except for the last two and classification layers are kept non-trainable. To avoid updating the saved state of the parameters substantially, we train it with a small learning rate for four epochs, it increases in 1000 steps to  $1e - 4$  and then decays linearly. The gradients are also clipped as discussed by You et al. (2019) and the AdamW optimizer is used to optimize the model (Loshchilov and Hutter, 2017). For machine learning models, we examine the feature importance scores from the XGBoost model for feature selection (Zheng et al., 2017). We only consider features that have relative feature importance scores more than 0.30, i.e. word  $n$ -grams *tf-idf* scores for  $n = \{1, 2, 3\}$  and character  $k$ -gram *tf-idf* scores for  $k = 1$  to 6, and the character count of the spans. We use the logistic regression model with L2 penalty and class weights inversely proportional to their count in the training set to get the technique classification probabilities. The output logits from BERT’s classification layer are scaled (min-max scaling) and concatenated with the scaled probabilities obtained from the logistic regression model. We pass this concatenated vector to another logistic regression classifier with the same hyper-parameters to get the final predictions.

## 5 Results

Embedding	Span-level Metrics			Token-level Metrics			
	Precision	Recall	F1	Precision	Recall	F1	Accuracy
BERT-Token	0.319	0.374	0.344	0.601	0.333	0.428	0.885
RoBERTa-Token	0.358	0.375	0.366	0.62	0.233	0.339	0.883
BERT-MG	0.304	0.409	0.349	0.846	0.786	<b>0.815</b>	<b>0.952</b>
RoBERTa-MG*	<b>0.347</b>	<b>0.391</b>	<b>0.368</b>	<b>0.852</b>	0.779	0.814	<b>0.952</b>
XLNet-MG	0.304	0.413	0.35	0.82	<b>0.794</b>	0.807	0.949
GPT-2-MG	0.289	0.406	0.337	0.805	0.752	0.778	0.942

Table 1: Results for the Span Identification sub-task on the development set (\*Final submission)

<sup>3</sup>Flair. A state-of-the-art NLP framework. <https://github.com/flairNLP/flair>

<sup>4</sup>Transformers. <https://github.com/huggingface/transformers>

## 5.1 Sub-Task 1: Span Identification

For the span-identification sub-task, we consider token-only baselines with *BERT-Token* and *RoBERTa-Token* models. These models only perform token-level binary-classification task and do not consider the sentence-level and article-level granularities. We compare these baselines to our proposed multi-granularity model (*BERT-MG*, *RoBERTa-MG*, *XLNet-MG* and *GPT-2-MG*). Table 1 shows the Span-level metrics for the official development set and the token-level metrics on our token-level validation split. Our multi-granularity model consistently achieves more than 94% token-level accuracy across all the four types of embeddings as compared to the 88% token-level accuracy shown by the token-only baselines. The multi-granularity model also shows some improvement on the span-level F1 score as compared to the token-only baselines. We get the best token-level F1 scores (0.815) with BERT and RoBERTa embeddings. RoBERTa embeddings outperforms all the other embeddings for the span-level metrics. Therefore, we use RoBERTa embeddings for the final submission to the task leaderboard. The performance for this sub-task can possibly be improved further by using trainable embeddings by fine-tuning their respective transformer-models. Our system ranks 21<sup>st</sup> on the final test-set leaderboard with a span-level F1 score of 0.387.

Model	Validation Split		Development Set	
	Micro F1	Macro F1	Micro F1	Macro F1
Logistic Regression	0.54	0.42	0.51	0.32
LSTM	0.55	0.32	0.49	0.29
LSTM (Glove)	0.59	0.36	0.55	0.36
CNN-LSTM (Glove)	0.56	0.28	0.48	0.26
BERT	<b>0.66</b>	0.42	<b>0.58</b>	0.38
BERT + Logistic Regression*	0.65	<b>0.45</b>	<b>0.58</b>	<b>0.43</b>
RoBERTa	0.63	0.38	0.54	0.32

Table 2: Results for the Technique Classification sub-task on the development set and validation split (\*Final submission)

## 5.2 Sub-Task 2: Technique Classification

For the technique-classification sub-task, our system ranks 18<sup>th</sup> with a micro-F1 score of 0.54 on the final test-set. We have shown the performances of our models on the development set in Table 2. Since this is an imbalanced multi-class classification task, we also report our models’ macro F1 scores to test our proposed systems’ generalization across all the propaganda techniques. Macro averaging gives equal importance to all classes. Table 2 unveils that in both validation and development splits, the combined BERT and logistic regression outperforms all models with a macro-f1 score of 0.45 and 0.43, respectively.

To further inspect the performance of our models across particular techniques, we also calculate the class-wise F1 scores as shown in Table 3. Machine learning models (logistic regression) achieve a relatively lower micro averaged F1 score of 0.514 on the development set, but give non-zero scores for all the classes. Stacked LSTM networks with randomly initialized embeddings do not achieve any improvement (0.490) on the development-set and conceive near-zero scores for multiple classes as well. This is due to the vast vocabulary of this relatively smaller dataset. Consequently, on using pre-trained GloVe embeddings with the stacked LSTM networks, we achieve a significant improvement (0.52).

Similarly, while CNN-LSTM models are relatively faster, they lower the overall performance (0.485) and conceive zero f1 scores for multiple classes. Transformer based models, namely BERT (0.578) and RoBERTa (0.550) outperform the above models significantly; however, they do not perform well for the minority classes. In fact, BERT gives zero or near zeros scores for some labels like *Appeal\_to\_Authority*; *Whataboutism*, *Straw\_Men*, *Red\_Herring*.

The ensemble of BERT with the logistic regression model does not improve the micro F1 score

	LR	LSTM	LSTM (Glove)	CNN-LSTM (Glove)	BERT	BERT-LR*	RoBERTa
Appeal_to_Authority	0.063	0.077	0.000	0.000	0.000	0.400	0.118
Appeal_to_fear-prejudice	0.160	0.205	0.292	0.321	0.376	0.372	0.317
Bandwagon, Reductio_ad_hitlerum	0.286	0.000	0.431	0.000	0.500	0.484	0.000
Black-and-White_Fallacy	0.167	0.051	0.154	0.000	0.154	0.145	0.000
Causal_Oversimplification	0.341	0.341	0.358	0.143	0.250	0.258	0.320
Doubt	0.444	0.431	0.493	0.497	0.517	0.479	0.503
Exaggeration, Minimisation	0.371	0.313	0.443	0.417	0.418	0.377	0.447
Flag-Waving	0.679	0.671	0.714	0.663	0.728	0.735	0.745
Loaded_Language	0.712	0.662	0.703	0.691	0.725	0.720	0.699
Name_Calling, labelling	0.591	0.539	0.645	0.682	0.693	0.700	0.691
Repetition	0.350	0.198	0.251	0.099	0.264	0.412	0.0546
Slogans	0.190	0.400	0.485	0.070	0.523	0.501	0.419
Thought-terminating_Cliches	0.061	0.121	0.000	0.000	0.167	0.240	0.111
Whataboutism, Straw_Men, Red_Herring	0.100	0.053	0.059	0.067	0.049	0.292	0.0588
<b>micro F1</b>	0.514	0.490	0.554	0.485	0.578	<b>0.584</b>	0.550
<b>macro F1</b>	0.322	0.290	0.359	0.260	0.383	<b>0.436</b>	0.320

Table 3: Class-wise F1 scores on the development set (\*Final submission)

significantly (0.584)(1% gain) but unlike other models it gives non-zero F1 score for all the classes, improving the macro F1 score from 0.38 to 0.43 (13.16% gain). The F1 scores increase from 0.0 to 0.4 for *Appeal\_to\_Authority* and from 0.049 to 0.292 for *Whataboutism, Straw\_Men, Red\_Herring*. We use this ensemble model as our final submission due to its better generalization across all the propaganda technique classes. Our results reveal that minority classes such as *Appeal to authority, Bandwagon, Reductio ad Hitlerum, Black and white fallacy* and *Thought terminating cliches* are relatively difficult to predict.

## 6 Reproducibility

We make our code publicly available as GitHub repositories. The source code for the proposed multi-granularity knowledge sharing model for sub-task 1 can be accessed at (<https://github.com/rajaswa/semEval2020-task11>). The source code for the proposed BERT ensemble with linguistic features based logistic regression model for sub-task 2 can be accessed at (<https://github.com/someshsingh22/News-Propaganda-Detection>). Our models and acquired results are available for benchmarking, comparison, and reproducibility. We do not share the experimental dataset as per the task guidelines.

## 7 Conclusion

In this paper, we proposed systems for the task of span identification and multi-class imbalanced technique classification of propaganda spans in news articles. We analyzed the performance of various machine learning and deep learning-based architectures for these high granularity tasks. On the span-identification sub-task test set, our multi-granularity knowledge sharing model gives a span-level F1 score of 0.387. For the technique classification task, our ensemble of pre-trained transformer model with logistic regression gives a micro F1 score of 0.54. We further infer the effectiveness of incorporating linguistic features and achieve non-zero F1 scores for all techniques and 13.6% gain in the macro-F1 score. Our results also unveil the limitations/ineffectiveness of deep learning models to capture the minority-class techniques.

We plan to extend our work by using trainable transformer-model embeddings and improve the performance of the span-identification sub-task. The work can further be enhanced by adding more granularities for knowledge-sharing. The proposed knowledge-sharing model may also be used for various closely-related tasks such as fake news and hate speech detection, given application-specific appropriate



objective functions are defined across multiple granularities for these tasks.

## Acknowledgements

We would like to thank the reviewers for their timely and constructive feedback on our work and the shared-task organizers for the opportunity and task-related resources. We would also like to thank the SemEval-2020 workshop organizers for the timely follow-ups and the workshop organization.

## References

- Kartik Aggarwal and Anubhav Sadana. 2019. Nsit@ nlp4if-2019: Propaganda detection from news articles using transfer learning. *EMNLP-IJCNLP 2019*, page 143.
- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual string embeddings for sequence labeling. In *COLING 2018, 27th International Conference on Computational Linguistics*, pages 1638–1649.
- Vian Bakir, Eric Herring, David Miller, and Piers Robinson. 2019. Organized persuasive communication: A new conceptual framework for research on public relations, propaganda and promotional culture. *Critical Sociology*, 45(3):311–328.
- Alberto Barrón-Cedeno, Giovanni Da San Martino, Israa Jaradat, and Preslav Nakov. 2019. Propopy: A system to unmask propaganda in online news. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9847–9848.
- Peter Bourgonje, Julian Moreno Schneider, and Georg Rehm. 2017. From clickbait to fake news detection: an approach based on detecting the stance of headlines to articles. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*, pages 84–89.
- Giovanni Luca Ciampaglia, Prashant Shiralkar, Luis M Rocha, Johan Bollen, Filippo Menczer, and Alessandro Flammini. 2015. Computational fact checking from knowledge networks. *PLoS one*, 10(6).
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, 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 5636–5646, Hong Kong, China, November. Association for Computational Linguistics.
- 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*.
- 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, June. Association for Computational Linguistics.
- Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. 2010. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*, 11(Feb):625–660.
- F. A. Gers, J. Schmidhuber, and F. Cummins. 1999. Learning to forget: continual prediction with lstm. In *1999 Ninth International Conference on Artificial Neural Networks ICANN 99. (Conf. Publ. No. 470)*, volume 2, pages 850–855 vol.2.
- Robert A Hackett. 1984. Decline of a paradigm? bias and objectivity in news media studies. *Critical Studies in Media Communication*, 1(3):229–259.
- Stefan Helmstetter and Heiko Paulheim. 2018. Weakly supervised learning for fake news detection on twitter. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 274–277. IEEE.

- Akshay Jain and Amey Kasbe. 2018. Fake news detection. In *2018 IEEE International Students' Conference on Electrical, Electronics and Computer Science (SCEECS)*, pages 1–5. IEEE.
- Garth S Jowett and Victoria O'donnell. 2018. *Propaganda & persuasion*. Sage Publications.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- 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. *CoRR*, abs/1907.11692.
- Ilya Loshchilov and Frank Hutter. 2017. Fixing weight decay regularization in adam. *CoRR*, abs/1711.05101.
- Rada Mihalcea and Carlo Strapparava. 2009. The lie detector: Explorations in the automatic recognition of deceptive language. In *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*, pages 309–312. Association for Computational Linguistics.
- 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.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svetlana 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.
- S. Singhal, R. R. Shah, T. Chakraborty, P. Kumaraguru, and S. Satoh. 2019. Spofake: A multi-modal framework for fake news detection. In *2019 IEEE Fifth International Conference on Multimedia Big Data (BigMM)*, pages 39–47.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *CoRR*, abs/1906.08237.
- Yang You, Jing Li, Jonathan Hseu, Xiaodan Song, James Demmel, and Cho-Jui Hsieh. 2019. Reducing BERT pre-training time from 3 days to 76 minutes. *CoRR*, abs/1904.00962.
- Huiting Zheng, Jiabin Yuan, and Long Chen. 2017. Short-term load forecasting using emd-lstm neural networks with a xgboost algorithm for feature importance evaluation. *Energies*, 10(8):1168.
- Chunting Zhou, Chonglin Sun, Zhiyuan Liu, and Francis Lau. 2015. A c-lstm neural network for text classification. *arXiv preprint arXiv:1511.08630*.