

Findings of the NLP4IF-2019 Shared Task on Fine-Grained Propaganda Detection

Giovanni Da San Martino¹ Alberto Barrón-Cedeño² Preslav Nakov¹

¹ Qatar Computing Research Institute, HBKU, Qatar

² Università di Bologna, Forlì, Italy

{gmartino, pnakov}@hbku.edu.qa a.barron@unibo.it

Abstract

We present the shared task on Fine-Grained Propaganda Detection, which was organized as part of the NLP4IF workshop at EMNLP-IJCNLP 2019. There were two subtasks. FLC is a fragment-level task that asks for the identification of propagandist text fragments in a news article and also for the prediction of the specific propaganda technique used in each such fragment (18-way classification task). SLC is a sentence-level binary classification task asking to detect the sentences that contain propaganda. A total of 12 teams submitted systems for the FLC task, 25 teams did so for the SLC task, and 14 teams eventually submitted a system description paper. For both subtasks, most systems managed to beat the baseline by a sizable margin. The leaderboard and the data from the competition are available at <http://propaganda.qcri.org/nlp4if-shared-task/>.

1 Introduction

Propaganda aims at influencing people’s mindset with the purpose of advancing a specific agenda. In the Internet era, thanks to the mechanism of sharing in social networks, propaganda campaigns have the potential of reaching very large audiences (Glowacki et al., 2018; Muller, 2018; Tardáguila et al., 2018).

Propagandist news articles use specific techniques to convey their message, such as *whataboutism*, *red Herring*, and *name calling*, among many others (cf. Section 3). Whereas proving intent is not easy, we can analyse the language of a claim/article and look for the use of specific propaganda techniques. Going at this fine-grained level can yield more reliable systems and it also makes it possible to explain to the user why an article was judged as propagandist by an automatic system.

With this in mind, we organised the shared task on fine-grained propaganda detection at the NLP4IF@EMNLP-IJCNLP 2019 workshop. The task is based on a corpus of news articles annotated with an inventory of 18 propagandist techniques at the fragment level. We hope that the corpus would raise interest outside of the community of researchers studying propaganda. For example, the techniques related to fallacies and the ones relying on emotions might provide a novel setting for researchers interested in Argumentation and Sentiment Analysis.

2 Related Work

Propaganda has been tackled mostly at the article level. Rashkin et al. (2017) created a corpus of news articles labelled as propaganda, trusted, hoax, or satire. Barrón-Cedeño et al. (2019) experimented with a binarized version of that corpus: propaganda vs. the other three categories. Barrón-Cedeno et al. (2019) annotated a large binary corpus of propagandist vs. non-propagandist articles and proposed a feature-based system for discriminating between them. In all these cases, the labels were obtained using distant supervision, assuming that all articles from a given news outlet share the label of that outlet, which inevitably introduces noise (Horne et al., 2018).

A related field is that of computational argumentation which, among others, deals with some logical fallacies related to propaganda. Habernal et al. (2018b) presented a corpus of Web forum discussions with instances of *ad hominem* fallacy. Habernal et al. (2017, 2018a) introduced *Argotario*, a game to educate people to recognize and create fallacies, a by-product of which is a corpus with 1.3k arguments annotated with five fallacies such as *ad hominem*, *red herring* and *irrelevant authority*, which directly relate to propaganda.

Unlike (Habernal et al., 2017, 2018a,b), our corpus uses 18 techniques annotated on the same set of news articles. Moreover, our annotations aim at identifying the minimal fragments related to a technique instead of flagging entire arguments.

The most relevant related work is our own, which is published in parallel to this paper at EMNLP-IJCNLP 2019 (Da San Martino et al., 2019) and describes a corpus that is a subset of the one used for this shared task.

3 Propaganda Techniques

Propaganda uses psychological and rhetorical techniques to achieve its objective. Such techniques include the use of logical fallacies and appeal to emotions. For the shared task, we use 18 techniques that can be found in news articles and can be judged intrinsically, without the need to retrieve supporting information from external resources. We refer the reader to (Da San Martino et al., 2019) for more details on the propaganda techniques; below we report the list of techniques:

1. Loaded language. Using words/phrases with strong emotional implications (positive or negative) to influence an audience (Weston, 2018, p. 6).

2. Name calling or labeling. Labeling the object of the propaganda as something the target audience fears, hates, finds undesirable or otherwise loves or praises (Miller, 1939).

3. Repetition. Repeating the same message over and over again, so that the audience will eventually accept it (Torok, 2015; Miller, 1939).

4. Exaggeration or minimization. Either representing something in an excessive manner: making things larger, better, worse, or making something seem less important or smaller than it actually is (Jowett and O'Donnell, 2012, p. 303), e.g., saying that an insult was just a joke.

5. Doubt. Questioning the credibility of someone or something.

6. Appeal to fear/prejudice. Seeking to build support for an idea by instilling anxiety and/or panic in the population towards an alternative, possibly based on preconceived judgments.

7. Flag-waving. Playing on strong national feeling (or with respect to a group, e.g., race, gender, political preference) to justify or promote an action or idea (Hobbs and Mcgee, 2008).

8. Causal oversimplification. Assuming one cause when there are multiple causes behind an issue. We include *scapegoating* as well: the transfer of the blame to one person or group of people without investigating the complexities of an issue.

9. Slogans. A brief and striking phrase that may include labeling and stereotyping. Slogans tend to act as emotional appeals (Dan, 2015).

10. Appeal to authority. Stating that a claim is true simply because a valid authority/expert on the issue supports it, without any other supporting evidence (Goodwin, 2011). We include the special case where the reference is not an authority/expert, although it is referred to as *testimonial* in the literature (Jowett and O'Donnell, 2012, p. 237).

11. Black-and-white fallacy, dictatorship. Presenting two alternative options as the only possibilities, when in fact more possibilities exist (Torok, 2015). As an extreme case, telling the audience exactly what actions to take, eliminating any other possible choice (*dictatorship*).

12. Thought-terminating cliché. Words or phrases that discourage critical thought and meaningful discussion about a given topic. They are typically short and generic sentences that offer seemingly simple answers to complex questions or that distract attention away from other lines of thought (Hunter, 2015, p. 78).

13. Whataboutism. Discredit an opponent's position by charging them with hypocrisy without directly disproving their argument (Richter, 2017).

14. Reductio ad Hitlerum. Persuading an audience to disapprove an action or idea by suggesting that the idea is popular with groups hated in contempt by the target audience. It can refer to any person or concept with a negative connotation (Teninbaum, 2009).

15. Red herring. Introducing irrelevant material to the issue being discussed, so that everyone's attention is diverted away from the points made (Weston, 2018, p. 78). Those subjected to a red herring argument are led away from the issue that had been the focus of the discussion and urged to follow an observation or claim that may be associated with the original claim, but is not highly relevant to the issue in dispute (Teninbaum, 2009).

| | |
|---|--|
| | Stereotyping_name_calling_or_labeling |
| 1 | Manchin says Democrats acted like babies at the SOTU |
| 2 | Democrat West Virginia Sen. Joe Manchin says his colleagues' refusal to stand or applaud during President Donald Trump's State of the Union speech was disrespectful and a signal that |
| | Black-and-white_Fallacy |
| | the party is more concerned with obstruction than it is with progress. |
| | |
| | Loaded_language |
| 4 | In a glaring sign of just how stupid and petty things have become in Washington these days, Manchin was invited on Fox News Tuesday morning to discuss how he was one of the only Democrats in the chamber for the State of the Union speech |
| | Exaggeration |
| | Loaded_language |
| | not looking as though Trump killed his grandma. |
| | |
| 6 | As Manchin noted, many Democrats bolted as soon as Trump's speech ended in an apparent effort to signal |
| | Exaggeration |
| | they can't even stomach being in the same room as the president |

Figure 1: The beginning of an article with annotations.

16. Bandwagon. Attempting to persuade the target audience to join in and take the course of action because “everyone else is taking the same action” (Hobbs and Mcgee, 2008).

17. Obfuscation, intentional vagueness, confusion. Using deliberately unclear words, to let the audience have its own interpretation (Suprabandari, 2007; Weston, 2018, p. 8). For instance, when an unclear phrase with multiple possible meanings is used within the argument and, therefore, it does not really support the conclusion.

18. Straw man. When an opponent’s proposition is substituted with a similar one which is then refuted in place of the original (Walton, 1996).

4 Tasks

The shared task features two subtasks:

Fragment-Level Classification task (FLC). Given a news article, detect all spans of the text in which a propaganda technique is used. In addition, for each span the propaganda technique applied must be identified.

Sentence-Level Classification task (SLC). A sentence is considered propagandist if it contains at least one propagandist fragment. We then define a binary classification task in which, given a sentence, the correct label, either *propaganda* or *non-propaganda*, is to be predicted.

5 Data

The input for both tasks consists of news articles in free-text format, collected from 36 propagandist and 12 non-propagandist news outlets¹ and then annotated by professional annotators. More details about the data collection and the annotation, as well as statistics about the corpus can be found in (Da San Martino et al., 2019), where an earlier version of the corpus is described, which includes 450 news articles. We further annotated 47 additional articles for the purpose of the shared task using the same protocol and the same annotators.

The training, the development, and the test partitions of the corpus used for the shared task consist of 350, 61, and 86 articles and of 16,965, 2,235, and 3,526 sentences, respectively. Figure 1 shows an annotated example, which contains several propaganda techniques. For example, the fragment *babies* on line 1 is an instance of both *Name-Calling* and *Labeling*. Note that the fragment *not looking as though Trump killed his grandma* on line 4 is an instance of *Exaggeration_or_Minimisation* and it overlaps with the fragment *killed his grandma*, which is an instance of *Loaded-Language*.

Table 1 reports the total number of instances per technique and the percentage with respect to the total number of annotations, for the training and for the development sets.

¹We obtained the gold labels about whether a given news outlet was propagandistic from the Media Bias Fact Check website: <http://mediabiasfactcheck.com/>

| Technique | Train (%) | Dev (%) |
|-----------------------------|---------------|-------------|
| Appeal to Authority | 116 (1.92) | 50 (5.92) |
| Appeal to fear / prejudice | 239 (3.96) | 103 (12.19) |
| Bandwagon | 13 (0.22) | 3 (0.36) |
| Black and White Fallacy | 109 (1.80) | 17 (2.01) |
| Causal Oversimplification | 201 (3.33) | 22 (2.60) |
| Doubt | 490 (8.11) | 39 (4.62) |
| Exaggeration, Minimisation | 479 (7.93) | 59 (6.98) |
| Flag Waving | 240 (3.97) | 63 (7.46) |
| Loaded Language | 2,115 (35.10) | 229 (27.10) |
| Name Calling, Labeling | 1,085 (17.96) | 87 (10.30) |
| Obfuscation, Intentional | | |
| Vagueness, Confusion | 11 (0.18) | 5 (0.59) |
| Red Herring | 33 (0.55) | 10 (1.18) |
| Reductio ad hitlerum | 54 (0.89) | 9 (1.07) |
| Repetition | 571 (9.45) | 101 (11.95) |
| Slogans | 136 (2.25) | 26 (3.08) |
| Straw Men | 13 (0.22) | 2 (0.24) |
| Thought-terminating Cliches | 79 (1.31) | 10 (1.18) |
| Whataboutism | 57 (0.94) | 10 (1.18) |

Table 1: Statistics about the gold annotations for the training and the development sets.

6 Setup

The shared task had two phases: In the development phase, the participants were provided labeled training and development datasets; in the testing phase, testing input was further provided.

Phase 1. The participants tried to achieve the best performance on the development set. A live leaderboard kept track of the submissions.

Phase 2. The test set was released and the participants had few days to make final predictions.

In phase 2, no immediate feedback on the submissions was provided. The winner was determined based on the performance on the test set.

7 Evaluation

FLC task. FLC is a composition of two sub-tasks: the identification of the propagandist text fragments and the identification of the techniques used (18-way classification task). While F_1 measure is appropriate for a multi-class classification task, we modified it to account for partial matching between the spans; see (Da San Martino et al., 2019) for more details. We further computed an F_1 value for each propaganda technique (not shown below for the sake of saving space, but available on the leaderboard).

SLC task. SLC is a binary classification task with imbalanced data. Therefore, the official evaluation measure for the task is the standard F_1 measure. We further report Precision and Recall.

8 Baselines

The baseline system for the SLC task is a very simple logistic regression classifier with default parameters, where we represent the input instances with a single feature: the length of the sentence. The performance of this baseline on the SLC task is shown in Tables 4 and 5.

The baseline for the FLC task generates spans and selects one of the 18 techniques randomly. The inefficacy of such a simple random baseline is illustrated in Tables 6 and 7.

9 Participants and Approaches

A total of 90 teams registered for the shared task, and 39 of them submitted predictions for a total of 3,065 submissions. For the FLC task, 21 teams made a total of 527 submissions, and for the SLC task, 35 teams made a total of 2,538 submissions.

Below, we give an overview of the approaches as described in the participants’ papers. Tables 2 and 3 offer a high-level summary.

9.1 Teams Participating in the Fragment-Level Classification Only

Team **newspeak** (Yoosuf and Yang, 2019) achieved the best results on the test set for the FLC task using 20-way word-level classification based on BERT (Devlin et al., 2019): a word could belong to one of the 18 propaganda techniques, to none of them, or to an auxiliary (token-derived) class. The team fed one sentence at a time in order to reduce the workload. In addition to experimenting with an out-of-the-box BERT, they also tried unsupervised fine-tuning both on the 1M news dataset and on Wikipedia. Their best model was based on the uncased base model of BERT, with 12 Transformer layers (Vaswani et al., 2017), and 110 million parameters. Moreover, oversampling of the least represented classes proved to be crucial for the final performance. Finally, careful analysis has shown that the model pays special attention to adjectives and adverbs.

Team **Stalin** (Ek and Ghanimifard, 2019) focused on data augmentation to address the relatively small size of the data for fine-tuning contextual embedding representations based on ELMo (Peters et al., 2018), BERT, and Grover (Zellers et al., 2019). The balancing of the embedding space was carried out by means of synthetic minority class over-sampling. Then, the learned representations were fed into an LSTM.

| Team | BERT | LSTM | Word Emb. | Char. Emb. | Features | Unsup. Tuning |
|--------------|------|------|-----------|------------|----------|---------------|
| CUNLP | | ✓ | ✓ | ✓ | | |
| Stalin | ✓ | ✓ | | | | |
| MIC-CIS | ✓ | ✓ | | | ✓ | |
| Ituorp | ✓ | | | | | |
| ProperGander | ✓ | ✓ | | | | |
| newspeak | ✓ | | | | | ✓ |

Table 2: Overview of the approaches for the fragment-level classification task.

| Team | BERT | LSTM | logreg | USE | CNN | Embeddings | Features | Context |
|--------------|------|------|--------|-----|-----|------------|----------|---------|
| NSIT | ✓ | ✓ | | | | | | |
| CUNLP | ✓ | | ✓ | | | | ✓ | |
| JUSTDeep | ✓ | ✓ | | | | ✓ | ✓ | |
| Tha3aroon | ✓ | | | ✓ | | | | |
| LIACC | | ✓ | | | | ✓ | ✓ | |
| MIC-CIS | ✓ | | ✓ | | ✓ | ✓ | ✓ | |
| CAUnLP | ✓ | | | | | | | ✓ |
| YMJA | ✓ | | | | | | | |
| jinfen | ✓ | | ✓ | | | | ✓ | |
| ProperGander | ✓ | | | | | | | |

Table 3: Overview of the approaches used for the sentence-level classification task.

9.2 Teams Participating in the Sentence-Level Classification Only

Team **CAUnLP** (Hou and Chen, 2019) used two context-aware representations based on BERT. In the first representation, the target sentence is followed by the title of the article. In the second representation, the previous sentence is also added. They performed subsampling in order to deal with class imbalance, and experimented with BERT_{BASE} and BERT_{LARGE}.

Team **LIACC** (Ferreira Cruz et al., 2019) used hand-crafted features and pre-trained ELMo embeddings. They also observed a boost in performance when balancing the dataset by dropping some negative examples.

Team **JUSTDeep** (Al-Omari et al., 2019) used a combination of models and features, including word embeddings based on GloVe (Pennington et al., 2014) concatenated with vectors representing affection and lexical features. These were combined in an ensemble of supervised models: bi-LSTM, XGBoost, and variations of BERT.

Team **YMJA** (Hua, 2019) also based their approach on fine-tuned BERT. Inspired by *kaggle* competitions on sentiment analysis, they created an ensemble of models via cross-validation.

Team **jinfen** (Li et al., 2019) used a logistic regression model fed with a manifold of representations, including TF.IDF and BERT vectors, as well as vocabularies and readability measures.

Team **Tha3aroon** (Fadel and Al-Ayyoub, 2019) implemented an ensemble of three classifiers: two based on BERT and one based on a universal sentence encoder (Cer et al., 2018).

Team **NSIT** (Aggarwal and Sadana, 2019) explored three of the most popular transfer learning models: various versions of ELMo, BERT, and RoBERTa (Liu et al., 2019).

Team **Mindcoders** (Vlad et al., 2019) combined BERT, Bi-LSTM and Capsule networks (Sabour et al., 2017) into a single deep neural network and pre-trained the resulting network on corpora used for related tasks, e.g., emotion classification.

Finally, team **Ituorp** (Mapes et al., 2019) used an attention transformer using BERT trained on Wikipedia and BookCorpus.

9.3 Teams Participating in Both Tasks

Team **MIC-CIS** (Gupta et al., 2019) participated in both tasks. For the sentence-level classification, they used a voting ensemble including logistic regression, convolutional neural networks, and BERT, in all cases using FastText embeddings (Bojanowski et al., 2017) and pre-trained BERT models. Beside these representations, multiple features of readability, sentiment and emotions were considered. For the fragment-level task, they used a multi-task neural sequence tagger, based on LSTM-CRF (Huang et al., 2015), in conjunction with linguistic features. Finally, they applied sentence- and fragment-level models jointly.

| SLC Task: Test Set (Official Results) | | | | |
|---------------------------------------|------------------|----------------|-----------|--------|
| Rank | Team | F ₁ | Precision | Recall |
| 1 | ltuorp | 0.6323 | 0.6028 | 0.6648 |
| 2 | ProperGander | 0.6256 | 0.5649 | 0.7009 |
| 3 | YMJA | 0.6249 | 0.6252 | 0.6246 |
| 4 | MIC-CIS | 0.6230 | 0.5735 | 0.6818 |
| 5 | CUNLP | 0.6183 | 0.5778 | 0.6648 |
| 6 | Tha3aroon | 0.6138 | 0.5309 | 0.7274 |
| 7 | JUSTDeep | 0.6112 | 0.5792 | 0.6468 |
| 8 | CAUnLP | 0.6109 | 0.5180 | 0.7444 |
| 9 | LIPN | 0.5962 | 0.5241 | 0.6914 |
| 10 | LIACC | 0.5949 | 0.5090 | 0.7158 |
| 11 | aschern | 0.5923 | 0.6050 | 0.5800 |
| 12 | MindCoders | 0.5868 | 0.5995 | 0.5747 |
| 13 | jinfen | 0.5770 | 0.5059 | 0.6712 |
| 14 | guanggong | 0.5768 | 0.5039 | 0.6744 |
| 15 | Stano | 0.5619 | 0.6666 | 0.4856 |
| 16 | nlpseattle | 0.5610 | 0.6250 | 0.5090 |
| 17 | gw2018 | 0.5440 | 0.4333 | 0.7306 |
| 18 | SDS | 0.5171 | 0.6268 | 0.4400 |
| 19 | BananasInPajamas | 0.5080 | 0.5768 | 0.4538 |
| 20 | Baseline | 0.4347 | 0.3880 | 0.4941 |
| 21 | NSIT | 0.4343 | 0.5000 | 0.3838 |
| 22 | Stalin | 0.4332 | 0.6696 | 0.3202 |
| 23 | Antiganda | 0.3967 | 0.6459 | 0.2863 |
| 24 | Debunkers | 0.2307 | 0.3994 | 0.1622 |
| 25 | SBnLP | 0.1831 | 0.2220 | 0.1558 |
| 26 | Sberiboba | 0.1167 | 0.5980 | 0.0646 |

Table 4: Official test results for the SLC task.

Team **CUNLP** (Alhindi et al., 2019) considered two approaches for the sentence-level task. The first approach was based on fine-tuning BERT. The second approach complemented the fine-tuned BERT approach by feeding its decision into a logistic regressor, together with features from the Linguistic Inquiry and Word Count (LIWC)² lexicon and punctuation-derived features. Similarly to Gupta et al. (2019), for the fragment-level problem they used a Bi-LSTM-CRF architecture, combining both character- and word-level embeddings.

Team **ProperGander** (Madabushi et al., 2019) also used BERT, but they paid special attention to the imbalance of the data, as well as to the differences between training and testing. They showed that augmenting the training data by oversampling yielded improvements when testing on data that is temporally far from the training (by increasing recall). In order to deal with the imbalance, they performed cost-sensitive classification, i.e., the errors on the smaller positive class were more costly. For the fragment-level classification, inspired by named entity recognition, they used a model based on BERT using Continuous Random Field stacked on top of an LSTM.

²<http://liwc.wpengine.com/>

| SLC Task: Development Set | | | | |
|---------------------------|------------------|----------------|-----------|--------|
| Rank | Team | F ₁ | Precision | Recall |
| 1 | Tha3aroon | 0.6883 | 0.6104 | 0.7889 |
| 2 | KS | 0.6799 | 0.5989 | 0.7861 |
| 3 | CAUnLP | 0.6794 | 0.5943 | 0.7929 |
| 4 | ProperGander | 0.6767 | 0.5774 | 0.8173 |
| 5 | JUSTDeep | 0.6745 | 0.6234 | 0.7347 |
| 6 | ltuorp | 0.6700 | 0.6351 | 0.7090 |
| 7 | CUNLP | 0.6649 | 0.6198 | 0.7171 |
| 8 | aschern | 0.6646 | 0.6104 | 0.7293 |
| 9 | jinfen | 0.6616 | 0.5800 | 0.7699 |
| 10 | YMJA | 0.6601 | 0.6338 | 0.6887 |
| 11 | SBnLP | 0.6548 | 0.5674 | 0.7740 |
| 12 | guanggong | 0.6510 | 0.5737 | 0.7523 |
| 13 | LIPN | 0.6484 | 0.5889 | 0.7212 |
| 14 | Stalin | 0.6377 | 0.5957 | 0.6860 |
| 15 | Stano | 0.6374 | 0.6561 | 0.6197 |
| 16 | BananasInPajamas | 0.6276 | 0.5204 | 0.7902 |
| 17 | Kloop | 0.6237 | 0.5846 | 0.6684 |
| 18 | nlpseattle | 0.6201 | 0.6332 | 0.6075 |
| 19 | gw2018 | 0.6038 | 0.5158 | 0.7280 |
| 20 | MindCoders | 0.5858 | 0.5264 | 0.6603 |
| 21 | NSIT | 0.5794 | 0.6614 | 0.5155 |
| 22 | Summer2019 | 0.5567 | 0.6724 | 0.4749 |
| 23 | Antiganda | 0.5490 | 0.6609 | 0.4695 |
| 24 | Cojo | 0.5472 | 0.6692 | 0.4627 |
| 25 | Baseline | 0.4734 | 0.4437 | 0.5074 |
| 26 | gudetama | 0.4734 | 0.4437 | 0.5074 |
| 27 | test | 0.4734 | 0.4437 | 0.5074 |
| 28 | Visionators | 0.4410 | 0.5909 | 0.3518 |
| 29 | MaLaHITJuniors | 0.3075 | 0.4694 | 0.2286 |

Table 5: Results for the SLC task on the development set at the end of phase 1 (see Section 6).

10 Evaluation Results

The results on the test set for the SLC task are shown in Table 4, while Table 5 presents the results on the development set at the end of phase 1 (cf. Section 6).³ The general decrease of the F₁ values between the development and the test set could indicate that systems tend to overfit on the development set. Indeed, the winning team **ltuorp** chose the parameters of their system both on the development set and on a subset of the training set in order to improve the robustness of their system.

Tables 6 and 7 report the results on the test and on the development sets for the FLC task. For this task, the results tend to be more stable across the two sets. Indeed, team **newspeak** managed to almost keep the same difference in performance with respect to team **Antiganda**. Note that team **MIC-CIS** managed to reach the third position despite never having submitted a run on the development set.

³Upon request from the participants, we reopened the submission system for the development set for both tasks after the end of phase 2; therefore, Tables 5 and 7 might not be up to date with respect to the online leaderboard.

| FLC Task: Test Set (Official Results) | | | | |
|---------------------------------------|------------------|----------------|-----------|--------|
| Rank | Team | F ₁ | Precision | Recall |
| 1 | newspeak | 0.2488 | 0.2862 | 0.2200 |
| 2 | Antiganda | 0.2267 | 0.2882 | 0.1868 |
| 3 | MIC-CIS | 0.1998 | 0.2234 | 0.1808 |
| 4 | Stalin | 0.1453 | 0.1920 | 0.1169 |
| 5 | CUNLP | 0.1311 | 0.3234 | 0.0822 |
| 6 | aschern | 0.1090 | 0.0715 | 0.2294 |
| 7 | ProperGander | 0.0989 | 0.0651 | 0.2056 |
| 8 | Sberiboba | 0.0450 | 0.2974 | 0.0243 |
| 9 | BananasInPajamas | 0.0095 | 0.0095 | 0.0095 |
| 10 | JUSTDeep | 0.0011 | 0.0155 | 0.0006 |
| 11 | Baseline | 0.0000 | 0.0116 | 0.0000 |
| 12 | MindCoders | 0.0000 | 0.0000 | 0.0000 |
| 13 | SU | 0.0000 | 0.0000 | 0.0000 |

Table 6: Official test results for the FLC task.

11 Conclusion and Further Work

We have described the NLP4IF@EMNLP-IJCNLP 2019 shared task on fine-grained propaganda identification. We received 25 and 12 submissions on the test set for the sentence-level classification and the fragment-level classification tasks, respectively. Overall, the sentence-level task was easier and most submitted systems managed to outperform the baseline. The fragment-level task proved to be much more challenging, with lower absolute scores, but most teams still managed to outperform the baseline.

We plan to make the schema and the dataset publicly available to be used beyond NLP4IF. We hope that the corpus would raise interest outside of the community of researchers studying propaganda: the techniques related to fallacies and the ones relying on emotions might provide a novel setting for researchers interested in Argumentation and Sentiment Analysis.

As a kind of advertisement, Task 11 at SemEval 2020⁴ is a follow up of this shared task. It features two complimentary tasks:

Task 1 Given a free-text article, identify the propagandist text spans.

Task 2 Given a text span already flagged as propagandist and its context, identify the specific propaganda technique it contains.

This setting would allow participants to focus their efforts on binary sequence labeling for Task 1 and on multi-class classification for Task 2.

⁴<http://propaganda.qcri.org/semEval2020-task11/>

| FLC Task: Development Set | | | | |
|---------------------------|--------------|----------------|-----------|--------|
| Rank | Team | F ₁ | Precision | Recall |
| 1 | newspeak | 0.2422 | 0.2893 | 0.2084 |
| 2 | Antiganda | 0.2165 | 0.2266 | 0.2072 |
| 3 | Stalin | 0.1687 | 0.2312 | 0.1328 |
| 4 | ProperGander | 0.1453 | 0.1163 | 0.1934 |
| 5 | KS | 0.1369 | 0.2912 | 0.0895 |
| 6 | CUNLP | 0.1222 | 0.3651 | 0.0734 |
| 7 | aschern | 0.1010 | 0.0684 | 0.1928 |
| 8 | gudetama | 0.0517 | 0.0313 | 0.1479 |
| 9 | AMT | 0.0265 | 0.2046 | 0.0142 |
| 10 | esi | 0.0222 | 0.0308 | 0.0173 |
| 11 | ltuorp | 0.0054 | 0.0036 | 0.0107 |
| 12 | Baseline | 0.0015 | 0.0136 | 0.0008 |
| 13 | CAUnLP | 0.0015 | 0.0136 | 0.0008 |
| 14 | JUSTDeep | 0.0010 | 0.0403 | 0.0005 |

Table 7: Results for FLC task on the development set. The values refer to the end of phase 1 (see section 6)

Acknowledgments

This research is part of the Propaganda Analysis Project,⁵ which is framed within the Tanbih project.⁶ The Tanbih project aims to limit the effect of “fake news”, propaganda, and media bias by making users aware of what they are reading, thus promoting media literacy and critical thinking, which is arguably the best way to address disinformation and “fake news.” The project is developed in collaboration between the Qatar Computing Research Institute (QCRI), HBKU and the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL).

The corpus for the task was annotated by A Data Pro,⁷ a company that performs high-quality manual annotations.

References

- Kartik Aggarwal and Anubhav Sadana. 2019. NSIT@NLP4IF-2019: Propaganda detection from news articles using transfer learning. In (Feldman et al., 2019).
- Hani Al-Omari, Malak Abdullah, Ola AlTiti, and Samira Shaikh. 2019. JUSTDeep at NLP4IF 2019 Task 1: Propaganda detection using ensemble deep learning models. In (Feldman et al., 2019).
- Tariq Alhindi, Jonas Pfeiffer, and Smaranda Muresan. 2019. Fine-tuned neural models for propaganda detection at the sentence and fragment levels. In (Feldman et al., 2019).

⁵<http://propaganda.qcri.org>

⁶<http://tanbih.qcri.org>

⁷<http://www.aiidatapro.com>

- Alberto Barrón-Cedeño, Giovanni Da San Martino, Israa Jaradat, and Preslav Nakov. 2019. Proppy: A system to unmask propaganda in online news. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, AAAI '19, pages 9847–9848, Honolulu, HI, USA.
- Alberto Barrón-Cedeno, Israa Jaradat, Giovanni Da San Martino, and Preslav Nakov. 2019. Proppy: Organizing the news based on their propagandistic content. *Information Processing & Management*, 56(5):1849–1864.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. [Universal sentence encoder](#). *CoRR*, abs/1803.11175.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news articles. 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 2019, Hong Kong, China.
- Lavinia Dan. 2015. Techniques for the Translation of Advertising Slogans. In *Proceedings of the International Conference Literature, Discourse and Multicultural Dialogue*, LDMD '15, pages 13–23, Mures, Romania.
- 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*, NAACL-HLT '19, pages 4171–4186, Minneapolis, MN, USA.
- Adam Ek and Mehdi Ghanimifard. 2019. Synthetic propaganda embeddings to train a linear projection. In [\(Feldman et al., 2019\)](#).
- Ibrahim Fadel, Ali Tuffaha and Mahmoud Al-Ayyoub. 2019. Pretrained ensemble learning for fine-grained propaganda detection. In [\(Feldman et al., 2019\)](#).
- Anna Feldman, Giovanni Da San Martino, Alberto Barrón-Cedeño, Chris Brew, Chris Leberknight, and Preslav Nakov, editors. 2019. *Proceedings of the 2019 Workshop on Natural Language Processing for Internet Freedom (NLP4IF): censorship, disinformation, and propaganda*. Hong Kong, China.
- André Ferreira Cruz, Gil Rocha, and Henrique Lopes Cardoso. 2019. On sentence representations for propaganda detection: From handcrafted features to word embeddings. In [\(Feldman et al., 2019\)](#).
- Monika Glowacki, Vidya Narayanan, Sam Maynard, Gustavo Hirsch, Bence Kollanyi, Lisa-Maria Neudert, Phil Howard, Thomas Lederer, and Vlad Barash. 2018. News and political information consumption in Mexico: Mapping the 2018 Mexican Presidential election on Twitter and Facebook. Technical Report COMPROM DATA MEMO 2018.2, Oxford University, Oxford, UK.
- Jean Goodwin. 2011. Accounting for the force of the appeal to authority. In *Proceedings of the 9th International Conference of the Ontario Society for the Study of Argumentation*, OSSA '11, pages 1–9, Ontario, Canada.
- Pankaj Gupta, Khushbu Saxena, Usama Yaseen, Thomas Runkler, and Hinrich Schutze. 2019. Neural architectures for fine-grained propaganda detection in news. In [\(Feldman et al., 2019\)](#).
- Ivan Habernal, Raffael Hannemann, Christian Polak, Christopher Klamm, Patrick Pauli, and Iryna Gurevych. 2017. Argotario: Computational argumentation meets serious games. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP '17, pages 7–12, Copenhagen, Denmark.
- Ivan Habernal, Patrick Pauli, and Iryna Gurevych. 2018a. Adapting serious game for fallacious argumentation to German: pitfalls, insights, and best practices. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation*, LREC '18, Miyazaki, Japan.
- Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. 2018b. Before name-calling: Dynamics and triggers of ad hominem fallacies in web argumentation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, NAACL-HLT '18, pages 386–396, New Orleans, LA, USA.
- Renee Hobbs and Sandra Mcgee. 2008. Teaching about propaganda: An examination of the historical roots of media literacy. *Journal of Media Literacy Education*, 6(62):56–67.
- 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 *Proceedings of the International AAAI Conference on Web and Social Media*, ICWSM '18, pages 518–527, Stanford, CA, USA.
- Wenjun Hou and Ying Chen. 2019. CAUnLP at NLP4IF 2019 shared task: Context-dependent BERT for sentence-level propaganda detection. In [\(Feldman et al., 2019\)](#).

- Yiqing Hua. 2019. Understanding BERT performance in propaganda analysis. In (Feldman et al., 2019).
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *CoRR*, abs/1508.01991.
- John Hunter. 2015. Brainwashing in a large group awareness training? The classical conditioning hypothesis of brainwashing. Master’s thesis, University of Kwazulu-Natal, Pietermaritzburg, South Africa.
- Garth S. Jowett and Victoria O’Donnell. 2012. What is propaganda, and how does it differ from persuasion? In *Propaganda & Persuasion*, chapter 1, pages 1–48. Sage Publishing.
- Jinfen Li, Zhihao Ye, and Lu Xiao. 2019. Detection of propaganda using logistic regression. In (Feldman et al., 2019).
- 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.
- Harish Tayyar Madabushi, Elena Kochkina, and Castelle Michael. 2019. Cost-sensitive BERT for generalisable sentence classification on imbalanced data. In (Feldman et al., 2019).
- Norman Mapes, Anna White, Radhika Medury, and Sumeet Dua. 2019. Divisive language and propaganda detection using multi-head attention transformers with deep learning BERT-based language models for binary classification. In (Feldman et al., 2019).
- Clyde R. Miller. 1939. The Techniques of Propaganda. From “How to Detect and Analyze Propaganda,” an address given at Town Hall. The Center for learning.
- Robert Muller. 2018. Internet Research Agency Indictment. <http://www.justice.gov/file/1035477/download>.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP ’14, pages 1532–1543, Doha, Qatar.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, NAACL-HLT ’2018, pages 2227–2237, New Orleans, LA, USA.
- 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*, EMNLP ’17, pages 2931–2937, Copenhagen, Denmark.
- Monika L Richter. 2017. The Kremlin’s platform for ‘useful idiots’ in the West: An overview of RT’s editorial strategy and evidence of impact. Technical report, Kremlin Watch.
- Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. 2017. Dynamic routing between capsules. In *Proceedings of the Annual Conference on Neural Information Processing Systems 2017*, NIPS ’17, pages 3856–3866, Long Beach, CA, USA.
- Francisca Niken Vitri Suprabandari. 2007. American propaganda in John Steinbeck’s *The Moon is Down*. Master’s thesis, Sanata Dharma University, Yogyakarta, Indonesia.
- Cristina Tardáguila, Fabrício Benevenuto, and Pablo Ortellado. 2018. Fake news is poisoning Brazilian politics. WhatsApp can stop it. <https://www.nytimes.com/2018/10/17/opinion/brazil-election-fake-news-whatsapp.html>.
- Gabriel H Teninbaum. 2009. Reductio ad Hitlerum: Trumping the judicial Nazi card. *Michigan State Law Review*, page 541.
- Robyn Torok. 2015. Symbiotic radicalisation strategies: Propaganda tools and neuro linguistic programming. In *Proceedings of the Australian Security and Intelligence Conference*, pages 58–65, Perth, Australia.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the Conference on Neural Information Processing Systems*, NIPS ’17, pages 5998–6008, Long Beach, CA, USA.
- George-Alexandru Vlad, Mircea-Adrian Tanase, Cristian Onose, and Dumitru-Clementin Cercel. 2019. Sentence-level propaganda detection in news articles with transfer learning and BERT-BiLSTM-Capsule model. In (Feldman et al., 2019).
- Douglas Walton. 1996. *The straw man fallacy*. Royal Netherlands Academy of Arts and Sciences.
- Anthony Weston. 2018. *A rulebook for arguments*. Hackett Publishing.
- Shehel Yoosuf and Yin Yang. 2019. Fine-grained propaganda detection with fine-tuned BERT. In (Feldman et al., 2019).
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. *CoRR*, abs/1905.12616.