

Leveraging Declarative Knowledge in Text and First-Order Logic for Fine-Grained Propaganda Detection

Ruize Wang^{1*}, Duyu Tang², Nan Duan², Wanjun Zhong³, Zhongyu Wei^{1†},
Xuanjing Huang¹, Daxin Jiang², Ming Zhou²

¹Fudan University, Shanghai, China

²Microsoft Corporation, Beijing, China

³Sun Yat-Sen University, Guangzhou, China

{rzwang18, zywei, xjhuang}@fudan.edu.cn

{zhongwj25}@mail2.sysu.edu.cn

{dutang, nanduan, djiang, mingzhou}@microsoft.com

Abstract

We study the detection of propagandistic text fragments in news articles. Instead of merely learning from input-output datapoints in training data, we introduce an approach to inject declarative knowledge of fine-grained propaganda techniques. Specifically, we leverage the declarative knowledge expressed in both first-order logic and natural language. The former refers to the logical consistency between coarse- and fine-grained predictions, which is used to regularize the training process with propositional Boolean expressions. The latter refers to the literal definition of each propaganda technique, which is utilized to get class representations for regularizing the model parameters. We conduct experiments on Propaganda Techniques Corpus, a large manually annotated dataset for fine-grained propaganda detection. Experiments show that our method achieves superior performance, demonstrating that leveraging declarative knowledge can help the model to make more accurate predictions.

1 Introduction

Propaganda is the approach deliberately designed with specific purposes to influence the opinions of readers. Different from the fake news which is entirely made-up and has no verifiable facts, propaganda is possibly built upon an element of truth, and conveys information with strong emotion or somewhat biased. This characteristic makes propaganda more effective and unnoticed through the rise of social media platforms. Some examples of propagandistic texts and definitions of corresponding techniques are shown in Figure 1.

We study the problem of fine-grained propaganda detection in this work, which is possible

*Work is done during internship at Microsoft Research Asia.

†Corresponding author.

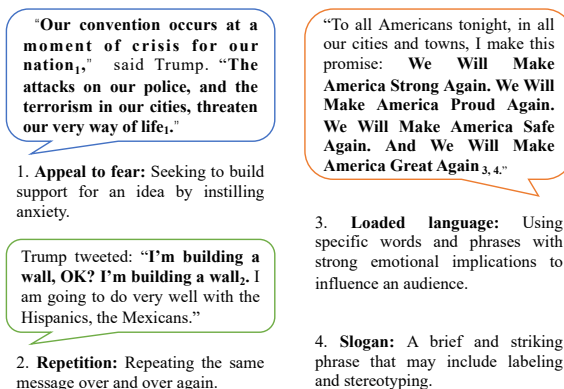


Figure 1: Examples of propagandistic texts, and definitions of corresponding propaganda techniques (**Bold** denotes propagandistic texts).

thanks to the recent release of Propaganda Techniques Corpus (Da San Martino et al., 2019). Different from earlier works (Rashkin et al., 2017; Wang, 2017) that mainly study propaganda detection at a coarse-grained level, namely predicting whether a document is propagandistic or not, the fine-grained propaganda detection requires to identify the tokens of particular propaganda techniques in news articles. Da San Martino et al. (2019) propose strong baselines in a multi-task learning manner, which are trained by binary detection of propaganda at sentence level and fine-grained propaganda detection over 18 techniques at token level. Such data-driven methods have the merits of convenient end-to-end learning and strong generalization, however, they cannot guarantee the consistency between sentence-level and token-level predictions. In addition, it is appealing to integrate human knowledge into data-driven approaches.

In this paper, we introduce an approach named *LatexPRO* that leverages logical and textual knowledge for **propaganda** detection. Following Da San Martino et al. (2019), we develop a BERT-based multi-task learning approach as the base

model, which makes predictions for 18 propaganda techniques at both sentence level and token level. Based on that, we inject two types of knowledge as additional objectives to regularize the learning process. Specifically, we exploit logic knowledge by transforming the consistency between sentence-level and token-level predictions with propositional Boolean expressions. Besides, we use the textual definition of propaganda techniques by first representing each of them as a contextual vector and then minimizing the distances to corresponding model parameters in semantic space.

We conduct extensive experiments on Propaganda Techniques Corpus (PTC) (Da San Martino et al., 2019), a large manually annotated dataset for fine-grained propaganda detection. Results show that our knowledge-augmented method significantly improves a strong multi-task learning approach. In particular, our model greatly improves precision, demonstrating leveraging declarative knowledge expressed in both first-order logic and natural language can help the model to make more accurate predictions. More importantly, further analysis indicates that augmenting the learning process with declarative knowledge reduces the percentage of inconsistency in model predictions.

The contributions of this paper are summarized as follows:

- We introduce an approach to leverage declarative knowledge expressed in both first-order logic and natural language for fine-grained propaganda techniques.
- We utilize both types of knowledge as regularizers in the learning process, which enables the model to make more consistent between sentence-level and token-level predictions.
- Extensive experiments on the PTC dataset (Da San Martino et al., 2019) demonstrate that our method achieves superior performance with high F_1 and precision.

2 Task

Task Definition. Following the previous work (Da San Martino et al., 2019), we conduct experiments on two different granularities tasks: sentence-level classification (SLC) and fragment-level classification (FLC). Formally, in both tasks, the input is a plain-text document d . A document includes a set of propagandistic fragments T , in

| Propaganda Technique | Instances | | |
|---------------------------------|--------------|------------|------------|
| | Train | Dev | Test |
| Loaded Language | 1,811 | 127 | 177 |
| Name Calling,Labeling | 931 | 68 | 86 |
| Repetition | 456 | 35 | 80 |
| Doubt | 423 | 23 | 44 |
| Exaggeration,Minimisation | 398 | 37 | 44 |
| Flag-Waving | 206 | 13 | 21 |
| Appeal to fear-prejudice | 187 | 32 | 20 |
| Causal Oversimplification | 170 | 24 | 7 |
| Slogans | 120 | 3 | 13 |
| Black-and-White Fallacy | 97 | 4 | 8 |
| Appeal to Authority | 91 | 2 | 23 |
| Thought-terminating Cliches | 70 | 4 | 5 |
| Whataboutism | 55 | 1 | 1 |
| Reductio ad hitlerum | 44 | 5 | 5 |
| Red Herring | 24 | 0 | 9 |
| Straw Men | 11 | 0 | 2 |
| Obfus.,Int. Vagueness,Confusion | 10 | 0 | 1 |
| Bandwagon | 10 | 2 | 1 |
| Total | 5,114 | 380 | 547 |

Table 1: The statistics of all 18 propaganda techniques.

that each fragment is represented as a sequence of contiguous characters $t = [t_i, \dots, t_j] \subseteq d$. For the SLC task, the target is to predict whether a sentence is propagandistic which can be regarded as a binary classification problem. For the FLC task, the target is to predict a set S with propagandistic fragments $s = [s_m, \dots, s_n] \subseteq d$ and identify $s \in S$ to one of the propagandistic techniques.

Dataset. This paper utilizes Propaganda Techniques Corpus (PTC) (Da San Martino et al., 2019) for experiments. PTC is a manually annotated dataset for fine-grained propaganda detection, containing 293/ 57/ 101 articles and 14,857/ 2,108/ 4,265 corresponding sentences for training, validation and testing, respectively. Each article is annotated with the start and end of the propaganda text span as well as the type of propaganda technique. As the annotations of the official testing set are not publicly available, we divided the official validation set into a validation set of 22 articles and a test set of 35 articles. The statistics of all 18 propaganda techniques and their frequencies (instances per technique) are shown as Table 1.

Evaluation. For the SLC task, we evaluate the models with precision, recall and micro-averaged F_1 scores. As for the FLC task, we adopt the evaluation script provided by Da San Martino et al. (2019) to calculate precision, recall, and micro-

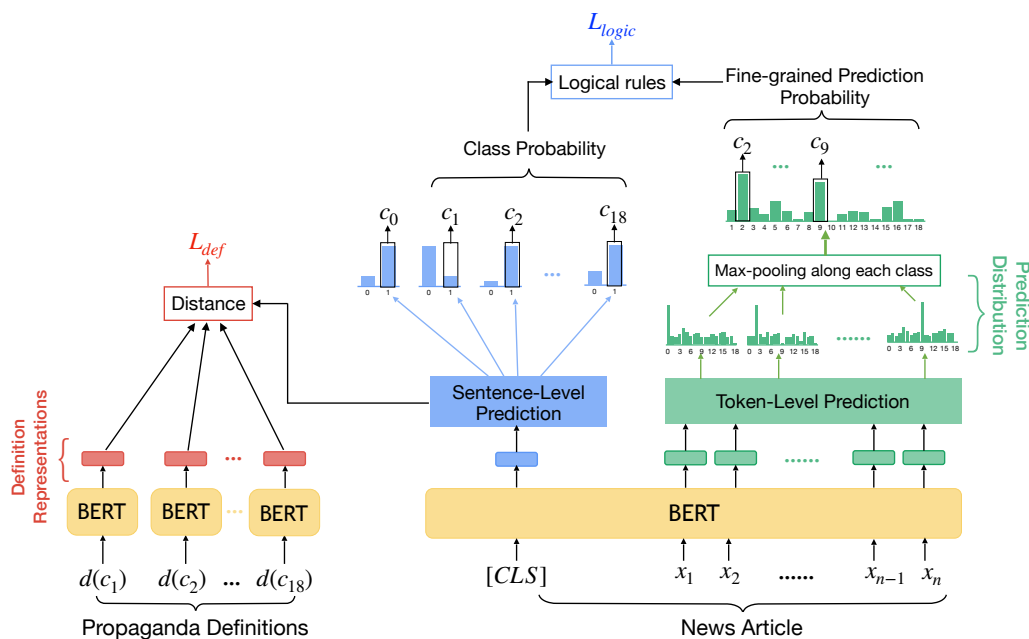


Figure 2: Overview of our proposed model. A BERT-based multi-task learning approach is adopted to make predictions for 18 propaganda techniques at both sentence level and token level. We introduce two types of knowledge as additional objectives: (1) logical knowledge about the consistency between sentence-level and token-level predictions, and (2) textual knowledge from literal definitions of propaganda techniques.

averaged F_1 , in that giving partial credit to imperfect matches at the character level. The FLC task is evaluated on two kinds of measures: (1) *Full task* is the overall task, which includes detecting the existence of propaganda techniques in text fragments and identifying the type of them, while (2) *Spans* is a special case of the *Full task*, which only considers the spans of propagandistic fragments except for their propaganda techniques.

3 Method

In this section, we present our approach *LatexPRO* as shown in Figure 2, which injects declarative knowledge of fine-grained propaganda techniques into neural networks. We first present our base model (§3.1), which is a multi-task learning framework that slightly extends the model of Da San Martino et al. (2019). Afterward, we introduce two ways to regularize the learning process with logical knowledge about the consistency between sentence-level and token-level predictions (§3.2) and textual knowledge from literal definitions of propaganda techniques (§3.3). At last, we describe the training and inference procedures (§3.4).

3.1 Base Model

To better exploit the sentence-level information and further benefit token-level prediction, we de-

velop a fine-grained multi-task method as our base model, which makes predictions for 18 propaganda techniques at both sentence level and token level. Inspired by the success of pre-trained language models on various natural language processing downstream tasks, we adopt BERT (Devlin et al., 2019) as the backbone model here. For each input sentence, the sequence is modified as “[CLS]sentence tokens[SEP]”. Specifically, on top of BERT, we add 19 binary classifiers for fine-grained sentence-level predictions, and one 19-way classifier for token-level predictions, where all classifiers are implemented as linear layers. At sentence level, we perform multiple binary classifications and this can further support leveraging declarative knowledge. The last representation of the special token [CLS] which is regarded as a summary of the semantic content of the input, is adopted to perform multiple binary classifications, including one binary classification to predict the existence of propaganda techniques, and 18 binary classifications to identify the types of them. We adopt sigmoid activation for each binary classifier. At token level, the last representation of each token is fed into a linear layer to predict the propaganda technique over 19 categories (i.e., 18 categories of propaganda techniques plus one category for “none of them”). We adopt Softmax activation for the

19-way classifier. Two different losses are applied for this multi-task learning process, including the sentence-level loss L_{sen} and the token-level loss L_{tok} . L_{sen} is the binary cross-entropy loss of multiple binary classifications. L_{tok} is the focal loss (Lin et al., 2017) of 19-way classification for each token, which could address the class imbalance problem.

3.2 Inject Logical Knowledge

There are some implicit logical constraints between sentence-level and token-level predictions. However, neural networks are less interpretable and need to be trained with a large amount of data to make it possible to learn such implicit logic. Therefore, we consider tackling the problems by exploiting logic knowledge. In particular, we propose to employ propositional Boolean expressions to explicitly regularize the model with a logic-driven objective, which improves the logical consistency between two different grained predictions, and makes our method more interpretable. For instance, in this work, if a propaganda class c is predicted by the multiple binary classifiers (indicates the sentence contains this propaganda technique), then the token-level predictions belonging to the propaganda class c should also exist. We thus consider the propositional rule $F = A \Rightarrow B$, formulated as:

$$\begin{aligned} P(F) &= P(A \Rightarrow B) \\ &= \neg P(A) \vee P(B) \\ &= 1 - P(A) + P(A)P(B) \\ &= P(A)(P(B) - 1) + 1 \end{aligned} \quad (1)$$

where A and B are two variables. Specifically, substituting $f_c(x)$ and $g_c(x)$ into above formula as $F = \forall c : f_c(x) \Rightarrow g_c(x)$, then the logic rule can be written as:

$$P(F) = P(f_c(x))(P(g_c(x)) - 1) + 1 \quad (2)$$

where x denotes the input, $f_c(x)$ is the binary classifier for the propaganda class c , and $g_c(x)$ is the probability of fine-grained predictions that contains x being category of c . $g_c(x)$ can be obtained by max-pooling over all the probability of predictions for class c . Note that the probabilities of the un-predicted class are set to 0 to prevent any violation, i.e., ensuring that each class has a probability corresponding to it. Our objective here is maximizing $P(F)$, i.e., minimizing $L_{logic} = -\log(P(F))$, to improve the logical consistency between coarse- and fine-grained predictions.

3.3 Inject Textual Knowledge

The literal definitions of propaganda techniques in this work, can be regarded as textual knowledge which contains useful semantic information. To exploit this kind of knowledge, we adopt an additional encoder to encode the literal definition of each propaganda technique. Specifically, for each definition, the input sequence is modified as “[CLS]definition[SEP]” and fed into BERT. We adopt the last representation of the special token [CLS] as each definition representation $D(c_i)$, where c_i represents the i -th propaganda technique. We calculate the Euclidean distance $dist_2$ between each predicted propaganda category representation $W(c_i)$ and the definition representation $D(c_i)$. Our objective is minimizing the textual definition loss L_{def} , which regularizes the model to refine the propaganda representations.

$$L_{def} = \sum_{i=1}^{18} dist_2(W(c_i), D(c_i)) \quad (3)$$

3.4 Training and Inference

Training. To train the whole model jointly, we introduce a weighted sum of losses L_j which consists of the token-level loss L_{tok} , fine-grained sentence-level loss L_{sen} , textual definition loss L_{def} and logical loss L_{logic} :

$$L_j = \alpha * L_{tok} + \beta * (L_{sen} + L_{def} * \lambda) + \gamma * L_{logic} \quad (4)$$

where hyper-parameters α , β , λ and γ are employed to control the tradeoff among losses. During the training stage, our goal is minimizing L_j .

Inference. For the SLC task, our method output the “propaganda” only if the probability of propagandistic binary classification for the positive class is above 0.7. This threshold is chosen according to the ratio of propaganda to non-propaganda samples in the training set. For the FLC task, to better exploit the coarse-grained (sentence-level) information to guide the fine-grained (token-level) prediction, we design a way that can explicitly make constraints on 19-way predictions when doing inference. Prediction probabilities of 18 fine-grained binary classifications above 0.9 are set to 1, and vice versa to 0. Then the Softmax probability of 19-way predictions (except for the “none of them” class) of each token is multiplied by the corresponding 18 probabilities of propaganda techniques. This means that our model is conservative, which makes

| Model | Spans | | | Full Task | | | |
|------------------------------------|--------------|--------------|----------------|--------------|--------------|----------------|----------------|
| | P | R | F ₁ | P | R | F ₁ | M _C |
| BERT (Da San Martino et al., 2019) | 50.39 | 46.09 | 48.15 | 27.92 | 27.27 | 27.60 | - |
| MGN (Da San Martino et al., 2019) | 51.16 | 47.27 | 49.14 | 30.10 | 29.37 | 29.73 | - |
| LatexPRO | 58.95 | 42.37 | 49.30 | 40.98 | 26.99 | 32.54 | 16.05 |
| LatexPRO (L) | 61.61 | 43.41 | 50.93 | 42.44 | 28.25 | 33.92 | 21.86 |
| LatexPRO (T) | 61.20 | 42.67 | 50.28 | 41.91 | 28.06 | 33.61 | 19.29 |
| LatexPRO (L+T) | 61.22 | 45.18 | 51.99 | 42.64 | 29.17 | 34.65 | 23.62 |

Table 2: Overall performance on fragment-level experiments (FLC task) in terms of Precision (P), recall (R) and F₁ scores on our test set. M_C denotes the metric of consistency between sentence-level predictions and token-level predictions. *Full task* is the overall task of detecting both propagandistic fragments and identifying the technique, while *Spans* is a special case of the *Full task*, which only considers the spans of fragments except for their propaganda techniques. Note that (L+T), (L), and (T) denote injecting of both logical and textual knowledge, only logical knowledge, and only textual knowledge, respectively.

predictions for the fragments of propaganda techniques only if with high confidence.

4 Experiments

4.1 Experimental Settings

In this paper, we conduct experiments on Propaganda Techniques Corpus (PTC)¹ (Da San Martino et al., 2019) which is a large manually annotated dataset for fine-grained propaganda detection, as detailed in Section 2. F₁ score is adopted as the final metric to represent the overall performance of models. We select the best model on the dev set.

We adopt BERT-base-cased (Devlin et al., 2019) as the pre-trained model. We implement our model using Huggingface (Wolf et al., 2019). We use AdamW as the optimizer. In our best model on the dev set, the hyper-parameters in loss optimization are set as $\alpha = 0.8$, $\beta = 0.2$, $\lambda = 0.001$ and $\gamma = 0.001$. We set the max sequence length to 256, the batch size to 16, the learning rate to 3e-5 and warmup steps to 500. We train our model for 20 epochs and adopt an early stopping strategy on the average validation F₁ score of *Spans* and *Full Task* with patience of 5 epochs. For all experiments, we set the random seed to 42 for reproducibility.

4.2 Models for Comparison

We compare our proposed methods with several baselines. Moreover, three variants of our method are provided to reveal the impact of each compo-

nent. The notations of LatexPRO (L+T), LatexPRO (L), and LatexPRO (T) denote our model which injects both logical and textual knowledge, only logical knowledge and only textual knowledge, respectively. Each of these models are described as follows.

BERT (Da San Martino et al., 2019) adds a linear layer on the top of BERT, and is fine-tuned on the SLC and FLC tasks, respectively.

MGN (Da San Martino et al., 2019) is a multi-task learning model, which regards the SLC task as the main task and drive the FLC task on the basis of the SLC task.

LatexPRO is our base model without leveraging any declarative knowledge.

LatexPRO (L) injects logical knowledge into *LatexPRO* by employing propositional Boolean expressions to explicitly regularize the model.

LatexPRO (T) arguments *LatexPRO* with textual knowledge in the literal definitions of propaganda techniques.

LatexPRO (L+T) is our full model in this paper.

4.3 Experiment Results and Analysis

Fragment-Level Propaganda Detection. The results for the FLC task are shown in Table 2. Our base model *LatexPRO* achieves better results than other baseline models, which verifies the effectiveness of our fine-grained multi-task learning structure. It is worth noting that, our full model *LatexPRO (L+T)* achieves superior boost than *MGN* by 10.06% precision and 2.85% F₁ on the *Spans* task, 12.54% precision and 4.92% F₁ on the *Full task*, which is considered as significant progress. This

¹Note that the annotations of the official PTC test set are not publicly available, thus we split the original dev set into dev and test set as Section 2. We use the released code (Da San Martino et al., 2019) to run the baseline.

| Propaganda Technique | MGN | | | LatexPRO | | | LatexPRO (L+T) | | |
|-----------------------------------|-------|-------|----------------|----------|-------|----------------|----------------|-------|----------------|
| | P | R | F ₁ | P | R | F ₁ | P | R | F ₁ |
| Appeal to Authority | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Appeal to fear-prejudice | 8.41 | 18.26 | 11.52 | 15.69 | 14.90 | 15.28 | 13.53 | 14.90 | 14.18 |
| Bandwagon | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Black-and-White Fallacy | 31.97 | 43.12 | 36.72 | 66.67 | 7.23 | 13.05 | 81.63 | 15.04 | 25.41 |
| Causal Oversimplification | 12.43 | 12.09 | 12.66 | 12.43 | 30.00 | 17.59 | 16.53 | 28.57 | 20.94 |
| Doubt | 27.12 | 12.38 | 17.00 | 18.06 | 9.09 | 12.09 | 40.82 | 9.26 | 15.10 |
| Exaggeration, Minimisation | 33.95 | 11.94 | 17.67 | 42.85 | 5.86 | 10.31 | 31.57 | 8.56 | 13.47 |
| Flag-Waving | 45.61 | 37.71 | 41.29 | 44.18 | 36.13 | 39.75 | 35.16 | 41.30 | 37.98 |
| Loaded Language | 37.20 | 46.45 | 41.31 | 51.69 | 39.19 | 44.58 | 50.28 | 44.39 | 47.15 |
| Name Calling, Labeling | 36.15 | 25.86 | 30.15 | 38.87 | 29.14 | 33.31 | 43.09 | 31.12 | 36.14 |
| Obfus., Int. Vagueness, Confusion | 0 | 0 | 0 | 100.00 | 98.61 | 99.30 | 50.00 | 98.61 | 66.35 |
| Red Herring | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Reductio ad hitlerum | 45.40 | 49.02 | 47.14 | 99.85 | 59.88 | 74.87 | 100.00 | 45.74 | 62.77 |
| Repetition | 35.05 | 24.09 | 26.93 | 46.06 | 28.75 | 35.40 | 48.24 | 26.86 | 34.51 |
| Slogans | 30.10 | 31.25 | 30.66 | 44.30 | 38.46 | 41.17 | 41.53 | 43.43 | 42.46 |
| Straw Men | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Thought-terminating Cliches | 21.05 | 23.85 | 22.36 | 90.83 | 14.80 | 25.45 | 89.49 | 19.60 | 32.16 |
| Whataboutism | 0 | 0 | 0 | 9.09 | 66.50 | 15.99 | 18.75 | 14.50 | 16.35 |

Table 3: Detailed performance on the full task of fragment-level experiments (FLC task) on our test set. Precision (P), recall (R) and F₁ scores per technique are provided.

demonstrates that leveraging declarative knowledge in text and first-order logic helps to predict the propaganda types more accurately. Moreover, our ablated models *LatexPRO (T)* and *LatexPRO (L)* both gain improvements over *LatexPRO*, while *LatexPRO (L)* gains more improvements than *LatexPRO (T)*. This indicates that injecting each kind of knowledge is useful, and the effect of different kinds of knowledge can be superimposed and uncoupled. It should be noted that, compared with baseline models, our models achieve a superior performance thanks to high precision, but the recall slightly loses. This is mainly because our models tend to make predictions for the high confident propaganda types.

To further understand the performance of models for the FLC task, we make a more detailed analysis of each propaganda technique. Table 3 shows detailed performance on the *Full* task. Our models achieve precision and F₁ improvements of almost all the classes over baseline model, and can also predict some low-frequency propaganda techniques, e.g., *Whataboutism* and *Obfus., Int.* This further demonstrates that our method can stress class imbalance problem, and make more accurate predictions.

Sentence-Level Propaganda Detection. Table 4 shows the performances of different models for SLC. The results indicate that our model achieves superior performances over other baseline mod-

| Model | P | R | F ₁ |
|------------------------------------|--------------|--------------|----------------|
| Random | 30.48 | 51.04 | 38.16 |
| All-Propaganda | 30.54 | 100.00 | 46.80 |
| BERT (Da San Martino et al., 2019) | 58.26 | 57.81 | 58.03 |
| MGN (Da San Martino et al., 2019) | 57.41 | 62.50 | 59.85 |
| LatexPRO | 56.18 | 69.79 | 62.25 |
| LatexPRO (L) | 56.53 | 73.17 | 63.79 |
| LatexPRO (T) | 58.33 | 67.50 | 62.58 |
| LatexPRO (L+T) | 59.04 | 71.66 | 64.74 |

Table 4: Results on sentence-level experiments (SLC task) in terms of Precision (P), recall (R) and F₁ scores on our test set. *Random* is a baseline which predicts randomly, and *All-Propaganda* is a baseline always predicts the propaganda class.

els. Compared with *MGN*, our *LatexPRO (L+T)* increases the precision by 1.63%, recall by 9.16% and F₁ score by 4.89%. This demonstrates the effectiveness of our model, and shows that our model can find more positive samples which will further benefit the token-level predictions for FLC.

4.4 Effectiveness of Improving Consistency

We further define the following metric M_C to measure the consistency between sentence-level predictions Y_c which is a set of predicted propaganda technique classes, and token-level predictions Y_t which is a set of predicted propaganda techniques for input tokens:

$$M_C(Y_c, Y_t) = \frac{1}{|Y_t|} \sum_{y_t \in Y_t} \mathbb{1}_{Y_c}(y_t) \quad (5)$$

| |
|---|
| Doc ID: 999000155 · Title: Arizona Democrat Senate Candidate Kyrsten Sinema Refuses To Retract Saying It’s OK For Americans To Join Taliban |
| Ground-truth: [...] “I go over there, and I’m fighting for the Taliban [...]” Sinema responded: “ Fine. I don’t care if you go and do that, go ahead. ” [...] The Left maintains that Islamic jihad terror is not a problem — it’s just a reaction to the evil deeds of the U.S. and Israel. [...]Sinema is proof: the Left hates America , and considers “right-wing extremists,” [...] It used to be that this fact was dismissed as hysterical hyperbole . |
| MGN: [...] “I go over there, and I’m fighting for the Taliban [...]” Sinema responded: “ Fine. I don’t care if you go and do that, go ahead. ” [...] The Left maintains that Islamic jihad terror is not a problem — it’s just a reaction to the evil deeds of the U.S. and Israel . [...]Sinema is proof: the Left hates America , and considers “right-wing extremists,” [...] It used to be that this fact was dismissed as hysterical hyperbole . |
| Our method: [...] “I go over there, and I’m fighting for the Taliban [...]” Sinema responded: “ Fine. I don’t care if you go and do that, go ahead. ” [...] The Left maintains that Islamic jihad terror is not a problem — it’s just a reaction to the evil deeds of the U.S. and Israel. [...]Sinema is proof: the Left hates America , and considers “right-wing extremists,” [...] It used to be that this fact was dismissed as hysterical hyperbole . |

Figure 3: Qualitative comparison of 2 different models on a news article. The baseline *MGN* predicts spans of fragments with wrong propaganda techniques, while our method can make more accurate predictions. Here are 5 propaganda techniques: 1.Thought-terminating Cliches, 2.Loaded Language, 3.Causal Oversimplification, 4.Flag waving and 5.Repetition. (Best viewed in color)

where $|Y_t|$ denotes a normalizing factor, $\mathbb{1}_A(x)$ represents the indicator function:

$$\mathbb{1}_A(x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{if } x \notin A \end{cases} \quad (6)$$

Table 2 presents the consistency scores M_C . The higher the score indicates the better consistency. Results illustrate that our methods with declarative knowledge can substantially outperform the base model *LatexPRO*. Compared to the base model, our declarative-knowledge-augmented methods enrich the source information by introducing textual knowledge from propaganda definitions, and logical knowledge from implicit logical rules between predictions, which enables the model to make more consistent predictions.

4.5 Case Study

Figure 3 gives a qualitative comparison example between *MGN* and our *LatexPRO (L+T)*. Different colors represent different propaganda techniques. The results show that although *MGN* could predict the spans of fragments correctly, it fails to identify their techniques to some extent. However,




Figure 4: Visualization of confusion matrix result of our *LatexPRO (L+T)*, where O represents the none of them class.

our method shows promising results on both spans and specific propaganda techniques, which further confirms that our method can make more accurate predictions.

4.6 Error Analysis

Although our model has achieved the best performance, it still some types of propaganda techniques are not identified, e.g., Appeal to Authority and Red Herring as shown in Table 3. To explore why our model *LatexPRO (L+T)* cannot predict for those propaganda techniques, we compute a confusion matrix for the *Full Task* of FLC task, and visualize the confusion matrix using a heatmap as shown in Figure 4. We find that most of the off-diagonal elements are in class O which represents none of them. This demonstrates most of the cases are wrongly classified into O. We think this is due to the imbalance of the propaganda and non-propaganda categories in the dataset. Similarly, Straw Men, Red Herring and Whataboutism are the relatively low frequency of classes. How to deal with the class imbalance still needs further exploration.

5 Related work

Our work relates to fake news detection and the injection of first-order logic into neural networks. We will describe related studies in these two directions.

Fake news detection draws growing attention as the spread of misinformation on social media becomes easier and leads to stronger influence. Various types of fake news detection problems are intro-

duced. For example, there are 4-way classification of news documents (Rashkin et al., 2017), and 6-way classification of short statements (Wang, 2017). There are also sentence-level fact checking problems with various genres of evidence, including natural language sentences from Wikipedia (Thorne et al., 2018), semi-structured tables (Chen et al., 2020), and images (Zlatkova et al., 2019; Nakamura et al., 2019). Our work studies propaganda detection, a fine-grained problem that requires token-level prediction over 18 fine-grained propaganda techniques. The release of a large manually annotated dataset (Da San Martino et al., 2019) makes the development of large neural models possible, and also triggers our work, which improves a standard multi-task learning approach by augmenting declarative knowledge expressed in both first-order logic and natural language.

Neural networks have the merits of convenient end-to-end training and good generalization, however, they typically need a lot of training data and are not interpretable. On the other hand, logic-based expert systems are interpretable and require less or no training data. It is appealing to leverage the advantages from both worlds. In NLP community, the injection of logic to neural network can be generally divided into two groups. Methods in the **first** group regularize neural network with logic-driven loss functions (Xu et al., 2018; Fischer et al., 2019; Li et al., 2019). For example, Rocktäschel et al. (2015) target on the problem of knowledge base completion. After extracting and annotating propositional logical rules about relations in knowledge graph, they ground these rules to facts from knowledge graph and add a differentiable training loss function. Kruszewski et al. (2015) map text to Boolean representations, and derive loss functions based on implication at Boolean level for entailment detection. Demeester et al. (2016) propose lifted regularization for knowledge base completion to improve the logical loss functions to be independent of the number of grounded instances and to further extend to unseen constants, The basic idea is that hypernyms have ordering relations and such relations correspond to component-wise comparison in semantic vector space. Hu et al. (2016) introduce a teacher-student model, where the teacher model is a rule-regularized neural network, whose predictions are used to teach the student model. Wang and Poon (2018) generalize virtual evidence (Pearl, 2014) to arbitrary potential

functions over inputs and outputs, and use deep probabilistic logic to integrate indirection supervision into neural networks. More recently, Asai and Hajishirzi (2020) regularize question answering systems with symmetric consistency and symmetric consistency. The former creates a symmetric question by replacing words with their antonyms in comparison question, while the latter is for causal reasoning questions through creating new examples when positive causal relationship between two cause-effect questions holds.

The **second** group is to incorporate logic-specific modules into the inference process (Yang et al., 2017; Dong et al., 2019). For example, Rocktäschel and Riedel (2017) target at the problem of knowledge base completion, and use neural unification modules to recursively construct model similar to the backward chaining algorithm of Prolog. Evans and Grefenstette (2018) develop a differentiable model of forward chaining inference, where weights represent a probability distribution over clauses. Li and Srikumar (2019) inject logic-driven neurons to existing neural networks by measuring the degree of the head being true measured by probabilistic soft logic (Kimmig et al., 2012). Our approach belongs to the first direction, and to the best of knowledge our work is the first one that augments neural network with logical knowledge for propaganda detection.

6 Conclusion

In this paper, we propose a fine-grained multi-task learning approach, which leverages declarative knowledge to detect propaganda techniques in news articles. Specifically, the declarative knowledge is expressed in both first-order logic and natural language, which are used as regularizers to obtain better propaganda representations and improve logical consistency between coarse- and fine-grained predictions, respectively. Extensive experiments on the PTC dataset demonstrate that our knowledge-augmented method achieves superior performance with more consistent between sentence-level and token-level predictions.

Acknowledgments

This work is partially supported by National Natural Science Foundation of China (No. 71991471), Science and Technology Commission of Shanghai Municipality Grant (No.20dz1200600, No.18DZ1201000, 17JC1420200).

References

- Akari Asai and Hannaneh Hajishirzi. 2020. Logic-guided data augmentation and regularization for consistent question answering. *arXiv preprint arXiv:2004.10157*.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2020. Tabfact: A large-scale dataset for table-based fact verification. In *ICLR*.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeno, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news article. In *EMNLP-IJCNLP*, pages 5640–5650.
- Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. 2016. Lifted rule injection for relation embeddings. In *EMNLP*, pages 1389–1399.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186.
- Honghua Dong, Jiayuan Mao, Tian Lin, Chong Wang, Lihong Li, and Denny Zhou. 2019. Neural logic machines. In *ICLR*.
- Richard Evans and Edward Grefenstette. 2018. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61:1–64.
- Marc Fischer, Mislav Balunovic, Dana Drachler-Cohen, Timon Gehr, Ce Zhang, and Martin Vechev. 2019. D12: Training and querying neural networks with logic. In *ICML*, pages 1931–1941.
- Zhiting Hu, Xuezhe Ma, Zhengzhong Liu, Eduard Hovy, and Eric Xing. 2016. Harnessing deep neural networks with logic rules. In *ACL*.
- Angelika Kimmig, Stephen Bach, Matthias Broecheler, Bert Huang, and Lise Getoor. 2012. A short introduction to probabilistic soft logic. In *Proceedings of the NIPS Workshop on Probabilistic Programming: Foundations and Applications*, pages 1–4.
- German Kruszewski, Denis Paperno, and Marco Baroni. 2015. Deriving boolean structures from distributional vectors. *Transactions of the Association for Computational Linguistics*, 3:375–388.
- Tao Li, Vivek Gupta, Maitrey Mehta, and Vivek Srikumar. 2019. A logic-driven framework for consistency of neural models. In *EMNLP-IJCNLP*, pages 3922–3933.
- Tao Li and Vivek Srikumar. 2019. Augmenting neural networks with first-order logic. In *ACL*, pages 292–302.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988.
- Kai Nakamura, Sharon Levy, and William Yang Wang. 2019. Fakeddit: A new multimodal benchmark dataset for fine-grained fake news detection. In *LREC*, pages 6149–6157.
- Judea Pearl. 2014. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Elsevier.
- 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 *EMNLP*, pages 2931–2937.
- Tim Rocktäschel and Sebastian Riedel. 2017. End-to-end differentiable proving. In *Advances in Neural Information Processing Systems*, pages 3788–3800.
- Tim Rocktäschel, Sameer Singh, and Sebastian Riedel. 2015. Injecting logical background knowledge into embeddings for relation extraction. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1119–1129.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: a large-scale dataset for fact extraction and verification. In *NAACL-HLT*, pages 809–819.
- Hai Wang and Hoifung Poon. 2018. Deep probabilistic logic: A unifying framework for indirect supervision. *arXiv preprint arXiv:1808.08485*.
- William Yang Wang. 2017. "liar, liar pants on fire": A new benchmark dataset for fake news detection. *arXiv preprint arXiv:1705.00648*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Van den Broeck. 2018. A semantic loss function for deep learning with symbolic knowledge. In *ICML*, pages 5498–5507.
- Fan Yang, Zhilin Yang, and William W Cohen. 2017. Differentiable learning of logical rules for knowledge base reasoning. In *Advances in Neural Information Processing Systems*, pages 2319–2328.
- Dimitrina Zlatkova, Preslav Nakov, and Ivan Koychev. 2019. Fact-checking meets fauxtography: Verifying claims about images. In *EMNLP-IJCNLP*, pages 2099–2108.