

Faking Fake News for Real Fake News Detection: Propaganda-loaded Training Data Generation

Kung-Hsiang Huang[♦] Kathleen McKeown[♦]

Preslav Nakov[◇] Yejin Choi[♡] Heng Ji[♦]

[♦] University of Illinois Urbana-Champaign [♦] Columbia University

[◇] Hamad Bin Khalifa University [♡] University of Washington [♦] Allen Institute for AI
{khhuang3, hengji}@illinois.edu

kathy@columbia.edu pnakov@hbku.edu.qa yejin@allenai.org

Abstract

While there has been a lot of research and many recent advances in neural fake news detection, defending against human-written disinformation remains underexplored. Upon analyzing current approaches for fake news generation and human-crafted articles, we found that there is a gap between them, which can explain the poor performance on detecting human-written fake news for detectors trained on automatically generated data. To address this issue, we propose a novel framework for generating articles closer to human-written ones. Specifically, we perform self-critical sequence training with natural language inference to ensure the validity of the generated articles. We then explicitly incorporate propaganda techniques into the generated articles to mimic how humans craft fake news. Eventually, we create a fake news detection training dataset, PROPANEWS, which includes 2,256 examples. Our experimental results show that detectors trained on PROPANEWS are 7.3% to 12.0% more accurate for detecting human-written disinformation than for counterparts trained on data generated by state-of-the-art approaches.¹

1 Introduction

The dissemination of falsified information can cause chaos, hatred, and trust issues among humans, and can eventually hinder the development of society (Dewatana and Adillah, 2021; Wasserman and Madrid-Morales, 2019). In particular, human-written disinformation², which is often used to manipulate certain populations, had catastrophic impact on multiple events, such as the 2016 US Presidential Election (Grinberg et al., 2019), Brexit (Bastos and Mercea, 2019), the COVID-19 pandemic (van Der Linden et al., 2020), and the

recent Russia’s assault on Ukraine. Hence, we are in urgent need of a defending mechanism against human-written disinformation.³

To construct such a mechanism, we need a substantial amount of training data to train the detectors. A naïve solution is to collect human-written news articles that contain inaccurate information by crawling untrustworthy news media and to recruit human annotators to label the veracity of each article. Unfortunately, news articles published by suspicious sources do not necessarily contain false information, which means that the annotators are required to perform fact-checking for every claim in each untrustworthy article. Additionally, these news outlets are often removed shortly after having been posted. Hence, such a solution is neither scalable nor reliable. A more appealing approach is to generate training data automatically that avoid these issues while enjoying the benefit of able to produce update-to-date disinformation.

Our goal here is to further advance the field of disinformation detection by generating a dataset that is closer to human-written disinformation, so that automatic detectors trained on it are more robust at detecting disinformation in human-written articles. We started by collecting human-written disinformative articles from untrustworthy sites⁴, and we analyzed around 40 of them that contain falsified information. Throughout our analysis, we found two characteristics of these human-written disinformation. First, about 33% of the articles use propaganda techniques to convince the audience that the fake information is authentic. These techniques usually involve the use of emotion-triggering terms, appealing to authorities or logical fallacies (Da San Martino et al., 2019) for increasing the credibility of the article. For instance,

¹The dataset and the code will be made publicly available.

²There are many types of *fake news*. We focus on text-only disinformation in this paper.

³WARNING: This paper contains disinformation that may be sensitive or offensive in nature.

⁴These new sources are rated *low* on factuality of reporting by mediabiasfactcheck.com

there is strong evidence showing that information is more persuasive and appears to be more true when it comes from sources perceived to be credible (Nadarevic et al., 2020), especially non-media sources such as authoritative individuals and domain experts (Nadarevic et al., 2020; Walter and Tukachinsky, 2020). Second, more than 55% of the articles that we analyzed contain inaccurate information mixed with correct information. Essentially, all claims, except for one or two, in these disinformation articles were factual. Thus, the fact that the majority of the claims in these articles are real makes these few false claims in these articles even more believable. Prior work has made great progress in generating fake news. For instance, Zellers et al. (2019) pre-trained a sequence-to-sequence (seq2seq) model, while Fung et al. (2021) presented a finer-grained generator that conditions on perturbed knowledge elements. A major issue with these approaches is that the vast majority of the contents in the articles they generate is inaccurate, which is in stark contrast with our observations about human-written disinformation, where accurate information accounts for a much greater proportion. Moreover, the articles generated using these methods do not explicitly use propaganda techniques, which are common in human-written disinformation (Da San Martino et al., 2019).

Given an authentic news article, we replace a salient sentence with a plausible but fake information using a sequence-to-sequence (seq2seq) language model. One issue is that if the generated sentence can be entailed by the replaced sentence, the generated sentence is actually accurate, which is not desirable. To penalize such an entailment, we fine tune the seq2seq model with a self-critical sequence training objective (Rennie et al., 2017) that uses a natural inference (NLI) model in addition to maximum likelihood estimation. Additionally, as a post-processing step, we further use the NLI model to filter out generated sentences that can be inferred from the replaced sentences. Then, we add propaganda techniques into the generated disinformation for mimicking how humans craft disinformation. In particular, we adopt two commonly used propaganda techniques: *appeal to authority* and *loaded language* (Da San Martino et al., 2019). The entire generation process is performed automatically to generate silver-standard training data for the detector. An example is shown in Table 1. For comparison, we recruited crowdsourcing workers

to correct grammatical errors and to validate that some of the generated texts are indeed fake to construct a gold-standard training dataset. Compared to the aforementioned annotation process, this data generation approach is much more viable as the annotators only need to fact-check the generated texts.

To evaluate the detector’s performance on human-written disinformation, we also collected 200 human-written articles (HUMANNEWS), with a balanced portion of real and fake articles. We compare our generation method with other state-of-the-art fake generation approaches, such as Zellers et al. (2019) and Shu et al. (2021). The results show that detectors are significantly better at detecting human-written disinformation when trained on our generated dataset. Our ablation studies further confirm the effectiveness of incorporating propaganda into the generated articles in identifying human-crafted disinformation.

Our contributions can be summarized as follows:

- We propose an effective method to automatically generate more realistic disinformation compared to prior work.
- We incorporate self-critical sequence training with an NLI model into the learning objectives of the generator to discourage generating sentences that can be entailed from the original context.
- We develop the first automatic methods for generating specific propaganda techniques such that the generated articles are closer to human-written disinformation.
- We demonstrate that detectors trained on our generated data, compared to generated articles using other methods, are better in detecting human-written disinformation.
- We release two disinformation detection datasets. (1) PROPANEWS: 2.2K articles generated using our approach and validated by humans. (2) HUMANNEWS: 200 human-written articles for evaluation in a real-world setting.

2 Training Data Generation

Our process of generating training data for propaganda-loaded disinformation consists of two main steps. First, we perform disinformation generation by replacing a salient sentence in an authentic

AJDABIYAH , Libya | Thu Apr 7 , 2011 6:34 pm EDT AJDABIYAH , Libya -LRB- Reuters -RRB- - Rebels fighting to overthrow Muammar Gaddafi said five of their fighters were killed when NATO planes mistakenly bombed a rebel tank column near the contested port of Brega in eastern Libya . In Washington , the head of U.S. Africa command told a Senate hearing the United States should not provide arms to the rebels without a better idea of who they were . Asked if there was an emerging stalemate in the seven-week-old conflict , General Carter Ham replied : “ I would agree with that at present , on the ground . ” In rebel-held eastern Libya , wounded rebels being brought to a hospital Ajdabiyah said their trucks and tanks were hit on Thursday by a NATO air strike outside Brega. NATO said it was investigating an attack by its aircraft on a tank column in the area along the Mediterranean coast on Thursday , saying the situation was “ unclear and fluid . ” Rebels said at least five of their fighters were killed when NATO planes mistakenly bombed a rebel tank column near the contested port. “ A number of vehicles were hit by a NATO strike ” , officers from UN concluded. The fighting for Brega , the only active front , has dragged on for a week and has entered a daily pattern of advances back and forth with neither side making major gains .

Table 1: An example of our generated fake news. Our approach first identifies a salient sentence in a given authentic news article (the ~~struckthrough~~ texts). Then, we generate a plausible but disinformative sentence that is coherent to the context (texts in orange). Finally, we generate propaganda to make it resemble human-written fake news (texts in blue)

news article with a plausible but inaccurate sentence (Section 2.1). Second, we incorporate propaganda techniques, including appeal to authority and loaded language, into the generated sentence (Section 2.2). Below, we describe each of these steps in detail.

2.1 Disinformation Generation

Our disinformation generation algorithm aims at two sub-goals: (i) replacing a salient sentence in the given article with a sequence of generated coherent texts that look plausible, and (ii) ensuring that the generated information cannot be entailed by the original masked-out sentence; otherwise, the generated texts will not be disinformative. To achieve the first sub-goal, we first identify salient sentences using an extractive summarization model, and we then perform mask infilling with BART (Lewis et al., 2020), a pre-trained seq2seq architecture that has demonstrated strong performance in multiple text generation tasks. The second sub-goal is accomplished using self-critical sequence training (Rennie et al., 2017; Bosselut et al., 2018) with a Natural Language Inference (NLI) component, which is used as a reward function for generation that meets this sub-goal.

Salient Sentence Identification A salient sentence is critical for the overall semantics of the article. When a salient sentence is manipulated or replaced, the news complex event described in the article may be drastically changed. Yet, there is no salient sentence identification dataset publicly available. Motivated by the fact that sentences included in an extractive summary are often of higher importance, we take the scores computed by an extractive summarization model (Liu and Lapata,

2019), which predicts how likely each sentence is to belong to the summary, to estimate the saliency of each sentence. Empirically, we found that this approach provides a reasonably good estimation for salient sentence identification. For each news outlet, we replace one sentence that has the highest extractive summarization score with our generated disinformation.

Mask Infilling with BART To perform infilling, we take an approach that is similar to that of Donahue et al. (2020), but we use BART (Lewis et al., 2020), a pre-trained language model with an encoder–decoder architecture. During training time, we randomly mask out a sentence y^* from a given article x . The bidirectional encoder first produces contextualized representations $\mathbf{h}_e = \text{Encoder}(\tilde{x})$ given the article with a masked-out sentence $\tilde{x} = x - y^*$. Then, the auto-regressive decoder learns a maximum likelihood estimation that aims to maximize the probability of generating the next token y_t^* at time step t given all tokens in previous time steps $\{y_0^*, \dots, y_{t-1}^*\}$ and the encoder hidden states \mathbf{h}_e by minimizing the negative log probability of generating y_t^* as follows:

$$\mathcal{L}_m = - \sum_{t=1}^T \log P(y_t^* | y_0^*, \dots, y_{t-1}^*, \mathbf{h}_e) \quad (1)$$

During inference time, rather than random masking, \tilde{x} is formed by masking out the sentence with the highest score computed by the extractive summarization model given the original document x , as discussed in the previous paragraph.

Self-critical Sequence Training BART optimized via maximum likelihood estimation alone is capable of generating coherent texts. However,

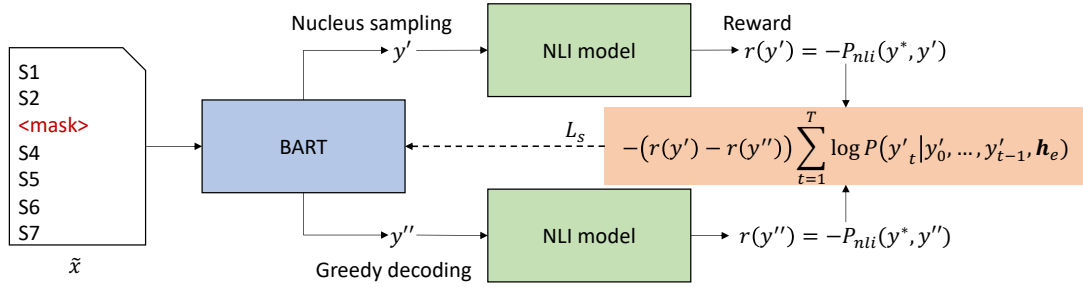


Figure 1: Illustration of our self-critical sequence training. Given a corrupted input article \tilde{x} , BART generates two sequences with Nucleus sampling and greedy decoding, respectively. The reward for each sequence is computed as the negative entailment probability $-P_{ent}$ as output from the NLI model.

although the generated texts y' may be very different from the originally masked out sentence y^* , there is no guarantee that y' contains incorrect information. If the generated texts y' can be entailed by the masked out sentence y^* , then y' is actually not disinformative. An example is shown in Figure 2. Here, except for the lack of details, the generated sentence y' delivers the same message as the masked out sentence y^* . To reduce the probability that y' can be entailed by y^* , we leverage self-critical sequence training (Rennie et al., 2017; Bosselut et al., 2018) that rewards the model for generating sequences that cannot be entailed by the masked-out sentences. Self-critical sequence

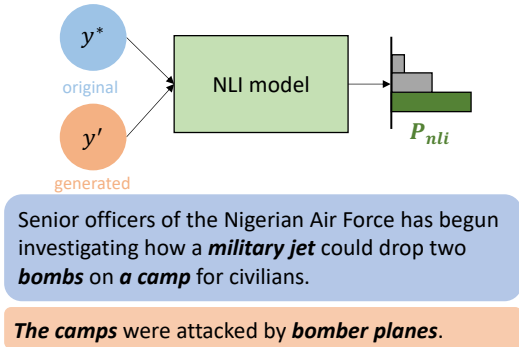


Figure 2: An example showing the NLI model predicts an entailment from the masked out sentence y^* to the generated sentence y' .

training (SCST) is a form of the REINFORCE algorithm (Williams, 1992) that allows for direct optimization on non-differentiable functions. Utilizing a baseline output y'' of the model itself to normalize the rewards, SCST avoids the challenge of directly estimating the reward signal or estimating normalization (Rennie et al., 2017). Since our goal is to avoid entailment from y^* to y' , we define the reward as the negative entailment probability produced by a ROBERTA-based (Liu et al., 2019) NLI

model fine-tuned on Multi-NLI (Williams et al., 2018)⁵,

$$r(y') = -P_{nli}(y^*, y'), \quad (2)$$

where $r(y')$ denotes the reward of the sequence sampled from the current policy y' , and $P_{nli}(y^*, y')$ is the probability that y^* entails y' . To generate y' , we use Nucleus Sampling (Holtzman et al., 2020) with $p = 0.96$ since such sampling method has demonstrated advantages over top-k sampling in *open-ended* generation (Holtzman et al., 2020; Zellers et al., 2019). Specifically, at each time step, we sample each token from the most probable set of words whose cumulative probability comprises at least p of the total vocabulary distribution. As for generating the baseline output y'' , we use greedy decoding by taking argmax at every timestep. Upon generating y' and y'' , we obtain the entailment probabilities of both sequences from the NLI model, and then compute the self-critical sequence training loss:

$$\mathcal{L}_s = -(r(y') - r(y'')) \sum_{t=1}^T \log P(y'_t | y'_0, \dots, y'_{t-1}, \mathbf{h}_e). \quad (3)$$

$r(y'')$ here can be viewed as a baseline reward, and $r(y') - r(y'')$ can be regarded as a normalized reward. This loss function encourages BART to generate y' when $r(y') > r(y'')$, whereas suppresses the probability of decoding y' when $r(y') < r(y'')$. An overview of SCST is shown in Figure 1.

The final objective function that BART learns to minimize is a weighted sum of Equation (1) and Equation (3),

⁵We use the fine-tuned NLI model from <https://huggingface.co/roberta-large-mnli>. Its accuracy is 90.2% on the dev set of MNLI, which is on par with state-of-the-art methods.

$$\mathcal{L}_{final} = \alpha\mathcal{L}_m + \beta\mathcal{L}_s, \quad (4)$$

where α and β are the weights for each loss⁶.

Post-processing To further ensure the quality of the disinformation generated, we reuse the NLI model discussed in the previous paragraph to filter out invalid outputs y^I that can be intailed from the masked-out sentence y^* , as demonstrated in Figure 2. We observe that the incorporation of the SCST loss (Equation (3)) into the training objective successfully reduces the invalid rate from 7.8% to 3.2%.

2.2 Propaganda Generation

After generating inaccurate information, we then incorporate propaganda into each generated article. In this work, we focus on generating two types of commonly used propaganda techniques, *loaded language* and *appeal to authority* (Da San Martino et al., 2019, 2020).

Appeal to Authority *Appeal to authority* is a propaganda technique that attempts to strengthen or invalidate an argument by referring to a statement made by authorities or experts (Da San Martino et al., 2019). These statements can be incorrect in two ways: (1) the statement itself is inaccurate, or (2) the authority never made such a statement. In this work, we focus on the latter case. For each article, we collect authority candidates Z from two sources, domain experts collected from Wikidata and entities gathered from the input article. With the Wikidata Query Service⁷, we collect experts from various domains, including economists, biologists, immunologists, and decision makers. Particularly, we specify the *occupation* (P108) of each expert and filter out entities that were born before 1940 to ensure recency. To consider only the impactful entities, we rank all the candidates based on the number of corresponding outgoing *statements* (i.e. connected concepts in Wikidata), inspired by PageRank (Page et al., 1999), and select the top 100 entities for each occupation into the candidate Z . Additionally, authorities that are more relevant to the local context are also considered. We include the person named entities extracted by a name tagger⁸ into the list of authority candidates based on

⁶Empirically, we set $\alpha = 1$ and $\beta = 0.01$.

⁷<https://query.wikidata.org/>

⁸<https://stanfordnlp.github.io/stanza/ner.html>

the our findings that more than 73% of the news articles contain authorities. More details about how authority candidates Z are collected can be found in Appendix B.

Once we collect a list of authority candidates Z , we can generate fake arguments made by each $z_i \in Z$ with the BART model that has already been fine-tuned in Section 2.1. In specific, a “<mask>” token is inserted right after the filled-in sentence y^I in the input article to BART so that it knows where to perform infilling. To inform BART that it should generate a statement made by one of the authorities, we prefix the decoder with a template such as [z_i confirmed that “”], where $z_i \in Z$ is the name of the authority. The prefix ends with an opening quotation mark to indicate that it should be followed by a statement made by authority z_i . Furthermore, to increase the diversity of the generated statements, we devise a variety of templates, as detailed in Appendix B. Finally, the best sequence s^* is selected with the lowest perplexity $s^* = \operatorname{argmin}_{s_i} \operatorname{Perplexity}(s_i)$ where s_i denotes the generated sequence using z_i as the authority.

Loaded Language *Loaded language* is another propaganda technique that utilizes emotion-triggering terms or phrases to influence audiences’ opinions (Da San Martino et al., 2019; Dimitrov et al., 2021). Often times, loaded language involves the use of sensational adverbs or adjectives to exaggerate a statement. Based on this observation, we utilize the propaganda dataset released by Da San Martino et al. (2019) where propaganda techniques are annotated at the fragment level (i.e. span level). The dataset contains 2,547 loaded language instances. Yet, not every instance contains adjectives or adverbs that are emotion-triggering. To create valid training data for loaded language generation, we first use SpaCy⁹ to perform part of speech tagging and dependency parsing, and then keep the examples where there exists an adverb pointing to a verb or an adjective pointing to a noun through dependency parsing edges. This results in 1,017 samples of valid loaded language instances. Some generated examples for *appeal to authority* and *loaded language* are shown in Table 2.

Upon collecting the training data for generating loaded language, we fine-tune another BART (Lewis et al., 2020) on this dataset. A naive method would be taking the article with emotion-triggering adverbs or adjectives removed as input to BART

⁹<https://spacy.io/>

Technique	Generated Disinformation and Propaganda
Appeal to Authority	Cairo’s Tahrir Square was the scene of clashes between protesters and police on Wednesday. “At least three people were killed and more than 600 were injured in the clashes,” said Egypt’s President.
Loaded Language	Cairo’s Tahrir Square was the scene of deadly clashes between protesters and police on Wednesday.

Table 2: Examples of the generated propaganda using two techniques, as shown by **texts in blue**. Given the same disinformation generated (i.e. the context), the first row demonstrates how loaded language is introduced by inserting an emotion-triggering term, while the second row shows how the argument is strengthened with reference to a statement made by authorities.

and using the original article as the decoding target. However, we found that around 25% of the time BART does not exactly reproduce the unmasked texts due to hallucination. This observation is consistent with Donahue et al. (2020)’s findings. To this end, we propose a two-step generation approach. First, we train BART to insert a “<mask>” token into the target sentence in the input document that is marked with special tokens. Then, BART learns to infill the “<mask>” with the similar approach as discussed in section 2.1 but without the SCST objective. Empirically, we found such approach successfully reduces the chance of failure in generating the exact unmasked contexts to around 2%.

2.3 Intermediate Pre-training

Directly fine-tuning BART on these two datasets does not yield satisfactory results as the size of TIMELINE17 (Tran et al., 2013) and the propaganda dataset (Da San Martino et al., 2019) are relatively small. To enhance the domain adaptation capabilities on small datasets, we perform intermediate pre-training (IPT) on the source news articles in the CNN/DM dataset (Hermann et al., 2015), a large summarization corpus containing more than 280K news articles from CNN and Daily Mail. The overall IPT objectives for disinformation generation and propaganda generation are mostly the same as described in Section 2.1 and Section 2.2, but with some minor changes due to different goals in the IPT phase. When performing IPT for disinformation generation, we remove \mathcal{L}_s from the final loss function (Equation (4)) as the goal for IPT is only to learn to generate coherent sentences. On the other hand, the objective for *loaded language* IPT is to enable BART to identify where to insert which adjectives or adverbs. Hence, we create training samples by gathering all appearances of adjectives pointing to a noun or adverbs pointing

to a verb via dependency parsing graphs.

3 The PROPANEWS Dataset

Using the generation method discussed in Section 2, we produce a set of realistic-looking disinformative articles for our proposed PROPANEWS dataset. To create a gold-standard training set for comparison, we hire human workers to validate that the generated articles are indeed disinformative. The details of this dataset are discussed in the following sections.

3.1 Data Source

When selecting the source of data to construct our dataset, we consider the following two criteria. First, the news articles must have high trustworthiness. This ensures that, except for our manipulated sentences, the rest of the articles are genuine. Second, the news events described in the articles must be important. If the disinformation presented in the articles is not important, such as gossip news, then these disinformative articles are not good subjects for studying disinformation detection. Motivated by these two criteria, we repurpose the TIMELINE17 dataset (Tran et al., 2013) as our source of data. TIMELINE17 contains 17 timelines, each of which corresponds to a news event. Each timeline is associated with a series of news articles that span across a wide time span, implying the high importance and impact of these news events. Additionally, the news articles are from trustworthy media, such as The New York Times and The Guardian. In total, there are 4,535 news articles in TIMELINE17.

3.2 Crowdsourcing

We use Amazon’s Mechanical Turk (AMT) to verify the quality and correctness of the generated disinformation. In total, there are around 400 unique crowdsourcing workers contributing to approxi-

mately 2,000 Human Intelligence Tasks (HITs). For each HIT, annotators are tasked to look for supporting evidence from trustworthy news media to determine whether the generated sentences are indeed *inaccurate*. Only those labeled as *inaccurate* will be included in PROPANEWS, while the *accurate* counterparts are discarded. Appendix C provides the details of the annotation interface.

To measure the inter-annotator agreement (IAA), we use the Worker Agreement With Aggregate (WAWA) score, following Ning et al. (2020) and Sheng et al. (2021)¹⁰. WAWA compares each annotator’s answer with the aggregated answer obtained via majority votes and micro-averages the results across all samples. The resulting WAWA precision, recall, and F_1 are 80.01%, 78.94%, and 79.47%, respectively, which indicates a moderate to high agreement.

4 Disinformation Detection

The disinformation detection task challenges detectors to determine whether a given input article contains inaccurate information or not. We experiment with four competitive detectors, including HDSF (Karimi and Tang, 2019), GROVER (Zellers et al., 2019), BERT (Devlin et al., 2019), ROBERTA (Liu et al., 2019). HDSF leverages the hierarchical structures of discourse-level features, such as dependency trees, to predict the veracity of a news article. GROVER is a unidirectional seq2seq model pre-trained on news-domain documents. It utilizes the encoder components of its generator counterpart for detection. The predictions are made by feeding the final hidden states of the encoder to a multi-layer perceptron. Similarly, BERT and ROBERTA take in the entire article as input and feed the representations of the first token to a classification head to determine the veracity of each article. In addition, INFOSURGEON is a more fine-grained detector that models the inconsistency at the level of knowledge elements across image and text modalities with a graph neural network-based architecture. Since we focus on text-only data, we set the input image features to zero vectors. All models are optimized using cross entropy. For fair comparison, we set the maximal sequence length to 512 and use the LARGE variants for all models. Details of the implementation can be found in

¹⁰We did not use other IAA metrics, such as Cohen’s Kappa (Cohen, 1960), as we expect vast majority of our generated disinformation to be inaccurate. WAWA provides a better approximation for inter-annotator agreement in our scenario.

Appendix D.

5 Experiments

In our experiments, we aim to (1) analyze the performance of different models on the PROPANEWS dataset, (2) examine the effect of various training data sets, and (3) investigate how much silver-standard data is equivalent to gold-standard data.

5.1 Data

PROPANEWS The PROPANEWS dataset consists of 2,256 distinct articles, with a balanced portion of fake and real documents. Within the fake articles, 30% of them use *appeal to authority*, another 30% include *loaded language*, and the remaining 40% simply contains inaccurate information. We split the data into 1,256: 500: 500 for training, validation, and testing.

Human-written news articles To evaluate the performance of our models on defending against human-written disinformation, we first collected around 60 disinformative articles debunked on politifact.com that were published between 2015 to 2020. Then, we further expanded the set of fake articles to 100 by fact checking news outlets published on untrustworthy news media¹. The fact checking procedure was done by a Computer Science majored graduate student. For real news articles, 100 articles from Los Angeles Times are curated. Eventually, this dataset (HUMANNEWS) contains 200 human-written news articles.

Other training data We compare the performance of detectors when trained on the following datasets to study the effectiveness of our dataset in defending against human-written disinformation. To understand the impact of human validation, we form the PN-SILVER dataset by resampling our generated articles but disregarding the annotator validation discussed in Section 3.2. In addition, we compare with two other approaches, GROVER-GEN (Zellers et al., 2019), FACTGEN (Shu et al., 2021) and FAKEEVENT, where fake articles in PN-SILVER are replaced with documents generated by according methods. GROVER-GEN generates headlines with condition on the original body texts, followed by body text generation conditioning on the generated headlines. FACTGEN enhances the factual consistency of the generated article with a fact retriever that fetches supporting information from external corpora. FAKEEVENT generates each sen-

Training Data	Detector	
	GROVER-LARGE	ROBERTA-LARGE
Without human validation (silver)		
GROVER-GEN	59.5 (± 3.0)	54.9 (± 2.6)
FAKEEVENT	44.7 (± 1.8)	42.4 (± 2.5)
FACTGEN	45.8 (± 7.7)	43.9 (± 3.2)
PN-SILVER	62.1 (± 4.6)	61.8 (± 1.3)
With human validation (gold)		
PROPANEWS	66.8 (± 2.0)	66.9 (± 2.1)
w/o AA	64.2 (± 2.1)	64.8 (± 2.9)
w/o LL	65.6 (± 2.9)	65.2 (± 1.8)
w/o AA & LL	63.8 (± 4.8)	63.8 (± 4.4)

Table 3: Accuracy (in %) on the HUMANNEWS dataset for different models when trained on various data sets. The bottom rows show different variants of PROPANEWS, with each component ablated. AA denotes *appeal to authority*, whereas LL refers to *loaded language*.

tence sequentially, with each generation step conditioning on the manipulated knowledge elements of the current sentence and previously generated sentences, following Fung et al. (2021). To ensure fair comparison, all generators take in the same set of authentic articles as inputs.

5.2 Results and Discussion

Human-written disinformation detection To study the effectiveness of detecting human-written disinformation, we train GROVER and ROBERTA on different training datasets and evaluate them on the HUMANNEWS dataset, as shown in Table 3. We see that both models perform the best when trained on PROPANEWS, compared to training on other datasets. Consider ablating human validation, detectors trained on PN-SILVER still outperform their counterparts trained on other datasets. This shows that our generative method produces articles that are more similar to human-written disinformation. To further verify such a finding, we measure the similarity between generated articles in different training data and the disinformative articles in the HUMANNEWS dataset using the MAUVE metric (Pillutla et al., 2021). MAUVE computes the similarity between two text distributions by summing the areas under a divergence curve, and it has been shown to produce better approximations than other metrics such as JS divergence (Martins et al., 2020). We found that the MAUVE score with HUMANNEWS for PROPANEWS and GROVER-GEN are 33.1% and 27.1%, respectively, suggesting that the generated documents in PROPANEWS are closer

to human-written disinformation. These results confirm that the advantage of our generated articles in defending against human-written disinformation is resulted from the closer gap between them.

In Table 4, we show two disinformative articles from HUMANNEWS where ROBERTA is able to classify them as inaccurate when trained on PN-SILVER but fails to do so when trained on GROVER-GEN. Both articles contain propaganda, which are incorporated into PN-SILVER but not into GROVER-GEN. This demonstrates that detectors are better at detecting human-written disinformation that has such properties.

Is propaganda generation helpful for disinformation detection? We have demonstrated that detectors trained on PROPANEWS perform better in identifying human-written disinformation compared to training on other datasets. We further conduct an ablation study to analyze the contributions of each propaganda technique. As shown in the bottom of Table 3, both *appeal to authority* and *loaded language* prove beneficial to enhancing models’ capabilities in detecting human-written disinformation. Furthermore, comparing PROPANEWS w/o AA& LL with other disinformation generation approaches, we find that both models trained on our generated data, even without the incorporation of propaganda techniques, still outperform their counterparts trained on other datasets. This illustrates that our generated articles, with disinformation mixed within real information, are more similar to human-written disinformation than articles generated using other approaches.

How important is human validation? The creation of gold-standard PROPANEWS involves human validation to ensure data quality. Although such a procedure is far more efficient than manual annotation of human-written disinformation, our ultimate goal is to completely relieve human efforts for even higher scalability. Therefore, we seek to understand the importance of such human validation process by comparing models trained on our generated data with (*gold*) or without (*silver*) human validation. In Table 5, we have shown the importance of human validation with the advantage of PROPANEWS over PN-SILVER. We further expand the experiment by training ROBERTA-LARGE on different scale of *silver* data, ranging from 1 time to 10 times the size of their *gold* counterpart (i.e. PROPANEWS). The 1 time *silver* data is equiva-

Article and Analysis

Article: ... “It is true that the Democratic Party should have put more resources into that election,” Sanders said on CNN’s “State of the Union” of the Thompson campaign. “But it is also true that he ran 20 points better than the Democratic candidate for president did in Kansas ...

Analysis: *Appealing to authority* is common in human-written fake news.

Article: ... **Fraudulent** White House Tells Businesses to Ignore Court Order on Vaccine Mandates Quick note : Tech giants are shutting us down ...

Analysis: The use of *loaded language* often indicates disinformation.

Table 4: Examples from HUMANNEWS where ROBERTA successfully predicts the veracity when trained on PN-SILVER, but classifies incorrectly when trained on GROVER-GEN.




Figure 3: Performance of ROBERTA on the HUMANNEWS dataset when trained on different size of data without human validation (*silver*), from 1 to 10 times of the data that involves human validation (*gold*). Each data point represents the average accuracy of four runs.

lent to the PN-SILVER dataset. Note that since the TIMELINE17 dataset only contains around 4K samples, we additionally crawl news articles from New York Times as inputs to our generator for the “5 times” to “10 times” experiments. The results are shown in Figure 3. We discover that when the size of *silver* data is increased to 5 times or above of the *gold* data size, the performance of the two models are comparable. The finding is likely due to the fact that even without human validation, the vast majority of the generated articles are indeed disinformative (~ 82%). This also reflects that when *silver*-standard data is abundant, human validation is not necessarily needed. However, our human validated *gold*-standard data can serve as a valuable resource for future research on pinpointing the disinformative passage in a given article as such directions may require higher-quality data.

How good is the generation quality? To evaluate the quality of our generation approach, we utilized AMT workers to rate the plausibility of 100 generated articles from PROPANEWS and de-

termine the degree by which their answer to this question is influenced by the generated propaganda. Each article is rated by 3 workers. As a comparison, we also ask AMT workers to rate the plausibility of 100 generated articles from GROVER-GEN. Details of the survey are discussed in Appendix E. The average plausibility scores for PROPANEWS and GROVER-GEN are 2.25 and 2.15 (out of 3), indicating that our generation approach has slight advantage over GROVER-GEN in terms of plausibility. Furthermore, among the articles in PROPANEWS that are rated highly plausible, 29.2% of the workers think that the generated propaganda highly affects their response (i.e. rated 3 out of 3) that the generated article is plausible. This demonstrates the effectiveness of our propaganda techniques in convincing the audience of the plausibility of generated articles.

Benchmarking detectors The performance of various detectors on the PROPANEWS dataset is shown in Table 5. We have the following observations. First, we find that ROBERTA and GROVER demonstrate advantages over BERT. This could be explained by the fact that ROBERTA and GROVER are pre-trained on news domain corpora, whereas BERT has no access to such domains during pre-training. In addition, we found that HDSF performs much worse than the other three models. This reflects that large-scale pre-training of language models brings more benefit to detection performance than explicit modeling of discourse-level features.

5.3 Remaining Challenges

To better understand the remaining disinformative articles that the detectors failed to identify, we conduct an analysis by comparing the ROBERTA predictions on the HUMANNEWS dataset and the la-

Detector	Dev Acc. (%)	Test Acc. (%)
HDSF	52.4 (± 0.6)	50.6 (± 2.4)
BERT	57.7 (± 1.0)	58.0 (± 1.2)
GROVER	60.3 (± 5.8)	63.3 (± 5.0)
ROBERTA	70.5 (± 0.3)	69.8 (± 1.1)

Table 5: Evaluation of various detectors on the PROPANEWS development and test set. We report the mean and standard deviation of four runs.

bels. Three major modeling capabilities required for successful detection are identified, as listed below:

Static knowledge enrichment About 30% of the misclassification is resulted from the lack of static knowledge that can be found in public databases, such as law dictionaries. For example, in this article¹¹, Alexandria Ocasio Cortez falsely states that the U.S. Immigration Customs Enforcement (ICE) is required to fill 34,000 beds every day. According to the Appropriations Act of 2016¹², however, ICE is only required to detain 34,000 available beds. Therefore, to detect such a misinformation, the detector needs to be enriched with static knowledge bases.

Dynamic knowledge acquisition Around 48% of the misclassified human-written disinformation are caused by the inability to acquire dynamic knowledge from new news sources. For instance, COVID-related articles are usually published after 2020, while ROBERTA was pre-trained on news articles released before 2019. It is very challenging for ROBERTA to detect disinformation of such topics unless the detector is equipped with the capabilities of acquiring dynamic knowledge from news articles. Particularly, ROBERTA achieves an accuracy of 69.0% on detecting fake articles published before 2019, but its accuracy drops to 51.9% when testing on articles published after 2019.

Multi-document reasoning The rest of the incorrect detection is caused by the lack of multi-document reasoning ability. For instance, a news article¹³ wrongly associates Hillary Clinton with a flawed immigration policy of the former government, and strengthens such a statement by referring to a Senate report and relevant news articles. However, the cited report does not mention Clinton, and

¹¹<https://bit.ly/3h5Qz00>

¹²<https://www.congress.gov/114/bills/hr2029/BILLS-114hr2029enr.pdf>

¹³<https://bit.ly/3t00V5v>

the other news articles contain disinformation. To correctly detect this piece of disinformation, detectors should be capable of reasoning across multiple documents.

6 Related Work

Fake News Generation and Detection Due to the rise of neural models and the potential threats of machine-generated fake news, prior study mostly focuses on how to automatically generate fake news with neural networks to defend against them. Zellers et al. (2019) pre-train a generator with the same architecture as GPT-2 (Radford et al., 2019) on a large-scale news corpus and demonstrate that this generator is effective in detecting neural fake news. More recently, Fung et al. (2021) improve the controllability of the generated fake news by conditioning the generator on knowledge elements, such as entities, relations and events, extracted from the original news article. Shu et al. (2021) enhances the factual consistency of the generated article by introducing a fact retriever that fetches relevant information from external corpora. Mosalanezhad et al. (2021) utilizes adversarial reinforcement learning to generate topic-preserving articles. These studies have developed methods for generating fake news that is hard to distinguish from real news to humans. Nevertheless, due to the overwhelming amount of inaccurate information introduced and the lack of propaganda techniques in the generated texts, these approaches are sub-optimal for detecting human-written fake news, as shown in Section 5.2. In contrast, our work generates fake news by incorporating propaganda techniques and preserving the majority of the correct information. Hence, our approach is more suitable for studying defense against human-written fake news. Additionally, since our released dataset is annotated with the exact offset of the disinformative passages, this work opens up future research opportunities on interpretable detection of fake news.

Propaganda Generation and Detection There is little previous study on propaganda generation. Zellers et al. (2019) is the only relevant work that we know of, which studies propaganda generation for communicating targeted disinformation. In contrast, our work focuses on generating specific propaganda techniques to make the generated articles closer to human-written fake news. To the best of our knowledge, we are the first to study the incorporation of specific propaganda techniques

into generated articles. Prior work on propaganda detection mostly focuses on document-level detection. Early works collect propaganda datasets using distant supervision (Rashkin et al., 2017; Barrón-Cedeno et al., 2019) by assigning the same propaganda label to each news outlet under the same source based on the news-media-level label of corresponding news source listed in trustworthy sites, such as US News & World Report. However, classifiers trained on such datasets may only learn to recognize the bias of each news source instead of propaganda (Da San Martino et al., 2020). Our dataset avoids such issues by explicitly incorporating propaganda into each generated article. Along another line of work, Da San Martino et al. (2019) present a fragment-level propaganda detection dataset, where specific propaganda techniques are labeled onto spans of text instead of each document. Recent approaches for detecting these propaganda techniques rely on pre-trained transformers (Morishita et al., 2020; Chernyavskiy et al., 2020; Feng et al., 2021; Tian et al., 2021). By contrast, we focus on detecting disinformative articles with propaganda signals.

7 Conclusions and Future Work

We have proposed a novel method for generating disinformation that are closer to human-written fake news. Evaluation on the human-written fake news dataset HUMANNEWS demonstrate the effectiveness of our generated data PROPANEWS in enabling better detection performance on human-written fake news. We hope that the datasets presented in this work, PROPANEWS and HUMANNEWS, can serve as enabling resources for human-written fake news detection and encourage future research along this direction. For future work, we plan to extend our approach to other languages and explore the cross-lingual detection settings. In addition, we are also interested in studying the generation of other propaganda techniques that involves logical fallacy, such as straw man and red herring, to enable more robust detection of human-written disinformation.

8 Ethical Statement and Broader Impact

Our objective for developing a generative approach that produces more realistic news articles is to advance the field of disinformation detection and to bring awareness that the current approaches for generating training data for fake news detection are

sub-optimal for defending against human-written fake news.

We acknowledge the dual-use concerns for such a generation framework, and therefore decided to release the codebase for only the detectors used in the experiments but not the generator.

Acknowledgement

This research is based upon work supported by U.S. DARPA SemaFor Program No. HR001120C0123. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Alberto Barrón-Cedeno, Israa Jaradat, Giovanni Da San Martino, and Preslav Nakov. 2019. Propopy: Organizing the news based on their propagandistic content. *Information Processing & Management*, 56(5):1849–1864.
- Marco T Bastos and Dan Mercea. 2019. The brexit botnet and user-generated hyperpartisan news. *Social science computer review*, 37(1):38–54.
- Antoine Bosselut, Asli Celikyilmaz, Xiaodong He, Jianfeng Gao, Po-Sen Huang, and Yejin Choi. 2018. Discourse-aware neural rewards for coherent text generation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 173–184, New Orleans, Louisiana. Association for Computational Linguistics.
- Anton Chernyavskiy, Dmitry Ilvovsky, and Preslav Nakov. 2020. Aschern at SemEval-2020 task 11: It takes three to tango: RoBERTa, CRF, and transfer learning. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1462–1468, Barcelona (online). International Committee for Computational Linguistics.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Giovanni Da San Martino, Stefano Cresci, Alberto Barrón-Cedeño, Seunghak Yu, Roberto Di Pietro, and Preslav Nakov. 2020. A survey on computational propaganda detection. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*.

- 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. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hernawan Dewatana and Siti Ummu Adillah. 2021. The effectiveness of criminal eradication on hoax information and fake news. *Law Development Journal*, 3(3):513–520.
- Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021. [Detecting propaganda techniques in memes](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6603–6617, Online. Association for Computational Linguistics.
- Chris Donahue, Mina Lee, and Percy Liang. 2020. [Enabling language models to fill in the blanks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2492–2501, Online. Association for Computational Linguistics.
- Zhida Feng, Jiji Tang, Jiayang Liu, Weichong Yin, Shikun Feng, Yu Sun, and Li Chen. 2021. [Alpha at SemEval-2021 task 6: Transformer based propaganda classification](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 99–104, Online. Association for Computational Linguistics.
- Yi Fung, Christopher Thomas, Revanth Gangi Reddy, Sandeep Polisetty, Heng Ji, Shih-Fu Chang, Kathleen McKeown, Mohit Bansal, and Avi Sil. 2021. [InfoSurgeon: Cross-media fine-grained information consistency checking for fake news detection](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1683–1698, Online. Association for Computational Linguistics.
- Nir Grinberg, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. 2019. Fake news on twitter during the 2016 us presidential election. *Science*, 363(6425):374–378.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *NIPS*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Hamid Karimi and Jiliang Tang. 2019. [Learning hierarchical discourse-level structure for fake news detection](#). 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 3432–3442, Minneapolis, Minnesota. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2019. [Text summarization with pretrained encoders](#). 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 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- 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](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Pedro Henrique Martins, Zita Marinho, and Andre F. T. Martins. 2020. [Sparse text generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4252–4273, Online. Association for Computational Linguistics.
- Terufumi Morishita, Gaku Morio, Hiroaki Ozaki, and Toshinori Miyoshi. 2020. [Hitachi at SemEval-2020](#)

- task 3: Exploring the representation spaces of transformers for human sense word similarity. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 286–291, Barcelona (online). International Committee for Computational Linguistics.
- Ahmadreza Mosallanezhad, Kai Shu, and Huan Liu. 2021. [Generating topic-preserving synthetic news](#). In *2021 IEEE International Conference on Big Data (Big Data)*, pages 490–499.
- Lena Nadarevic, Rolf Reber, Anne Josephine Helmecke, and Dilara Köse. 2020. Perceived truth of statements and simulated social media postings: an experimental investigation of source credibility, repeated exposure, and presentation format. *Cognitive Research: Principles and Implications*, 5(56).
- Qiang Ning, Hao Wu, Rujun Han, Nanyun Peng, Matt Gardner, and Dan Roth. 2020. [TORQUE: A reading comprehension dataset of temporal ordering questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1158–1172, Online. Association for Computational Linguistics.
- Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. [The pagerank citation ranking: Bringing order to the web](#). Technical Report 1999-66, Stanford InfoLab. Previous number = SIDL-WP-1999-0120.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. Mauve: Measuring the gap between neural text and human text using divergence frontiers. In *NeurIPS*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. [Truth of varying shades: Analyzing language in fake news and political fact-checking](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2931–2937, Copenhagen, Denmark. Association for Computational Linguistics.
- Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. [Self-critical sequence training for image captioning](#). *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Niloufar Salehi, Lilly C. Irani, Michael S. Bernstein, Ali Alkhatib, Eva Ogbe, Kristy Milland, and Clickhappier. 2015. [We are dynamo: Overcoming stalling and friction in collective action for crowd workers](#). In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, page 1621–1630, New York, NY, USA. Association for Computing Machinery.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2021. [“nice try, kiddo”: Investigating ad hominem in dialogue responses](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 750–767, Online. Association for Computational Linguistics.
- Kai Shu, Yichuan Li, Kaize Ding, and Huan Liu. 2021. [Fact-enhanced synthetic news generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13825–13833.
- Junfeng Tian, Min Gui, Chenliang Li, Ming Yan, and Wenming Xiao. 2021. [MinD at SemEval-2021 task 6: Propaganda detection using transfer learning and multimodal fusion](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 1082–1087, Online. Association for Computational Linguistics.
- Giang Binh Tran, Tuan Tran, Nam Khanh Tran, Mohammad Alrifai, and Nattiya Kanhabua. 2013. Leveraging learning to rank in an optimization framework for timeline summarization.
- Sander van Der Linden, Jon Roozenbeek, and Josh Compton. 2020. Inoculating against fake news about covid-19. *Frontiers in psychology*, 11:2928.
- Nathan Walter and Riva Tukachinsky. 2020. A meta-analytic examination of the continued influence of misinformation in the face of correction: How powerful is it, why does it happen, and how to stop it? *Communication Research*, 47:155–177.
- Herman Wasserman and Dani Madrid-Morales. 2019. An exploratory study of “fake news” and media trust in kenya, nigeria and south africa. *African Journalism Studies*, 40(1):107–123.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. In *Advances in Neural Information Processing Systems 32*.

A Qualitative Examples of Generated Articles

In Table 7, we show a comparison of generated articles given the same input data across different generative methods. Our approach produces articles with a small fraction of inaccurate information, which matches a property of human-written fake news discussed in Section 1.

B Appeal to Authority Details

To recap, we first gather a list of authorities Z for each article from Wikidata and the corresponding context. The best *appeal to authority* sequence s^* is selected with the lowest perplexity $s^* = \operatorname{argmin}_{s_i} \text{PPL}(s_i)$ where s_i denotes the generated sequence using z_i as the authority. However, this process results in every sequence s^* containing the the substring “confirms that”, which makes it trivial for detectors to classify these generated documents as fake by simply detecting such substrings. Therefore, we devise an algorithm to diversify the templates so that these generated articles are not easily detectable.

First, we define a set of verbs V that can be swapped with “confirms”: $V = \{\textit{said, concluded, confirmed, emphasized, stated, argued}\}$. Then, we diversify the generated structure of the generated sentence s^* by reordering the subject, verb, and object. Next, we swap the verb with a another verb from V . Finally, to diversify the context, we append a preposition from the preposition set $PP = \{\textit{on, at, in}\}$ to the output of the previous step, and then feed the sequence to BART to generate the context. An example of this process is provided in Table 6.

C Human Validation Details

In this section, we describe the details of human validation where AMT workers are tasked to validate whether the generated sentences contain inaccurate information. To ensure the annotation quality, only workers who have an acceptance rate greater than 95% and have more than 100 accepted HITs in the past are allowed to work on our annotation task. Each HIT was designed such that the annotators are rewarded \$12-\$15 per hour, which complies with the ethical research standards outlined by AMT (Salehi et al., 2015). In each HIT, the annotators are presented an article with the generated part marked in boldface. The questions and guidelines

are illustrated as follows. (Note that we only use the annotators’ response for Q1 to validate our generated data. The annotations for the other questions will be used for future research.)

Q1: Is the generated text in boldface **Accurate** or **Inaccurate**? (If you cannot find any supporting evidence, please select **Inaccurate**.) Note that: A statement (in quotation marks) made by a person is only accurate if this person actually made the exact same statement. If the statement in quotation marks is just a paraphrase of what the person actually said, then the statement is inaccurate.

- **Inaccurate:** Any false information presented in the generated text makes it inaccurate.
- **Accurate:** All the information in the generated text must be accurate.

Q2: Enter the URL of the news article you found that supports your decision in the previous response in the below box. Put down “from context” if the evidence can be found in the context.

Q3: Does the generated text in boldface delivers the same sentiment as the rest of the article?

- **False:** The sentiment of the generated text is NOT the same as the rest of the article.
- **True:** The sentiment of the generated text is the same as the rest of the article.

Q4: Is the discourse of the generated text in boldface consistent with the rest of the article?

- **False:** The discourse of the generated text is NOT consistent with the rest of the article.
- **True:** The discourse of the generated text is consistent with the rest of the article.

Q5: If there is any grammatical error or inconsistent discourse, please rewrite the correct the generated text and put it in the below box. Just put down the corrected generated text in bold is enough. For example, “Harry is a boy. He likes go to school.” Please put in “He likes to go to school.” in the box below.

D Detector Implementation Details

For BERT and ROBERTA experiments, we use AdamW (Loshchilov and Hutter, 2019) as the optimizer with a batch size of 2 and gradient accumulation steps of 8. We set the learning rate and

Step	Generated Sequence
1	Panmure Gordon analyst Peter Hitchens confirmed that “ the US government is likely to agree to reduce its estimate of the size of the spill, which would cut BP fines ”.
2	“ The US government is likely to agree to reduce its estimate of the size of the spill, which would cut BP fines, ” Panmure Gordon analyst Peter Hitchens confirmed.
3	“ The US government is likely to agree to reduce its estimate of the size of the spill, which would cut BP fines, ” Panmure Gordon analyst Peter Hitchens said.
4	“ The US government is likely to agree to reduce its estimate of the size of the spill, which would cut BP fines, ” Panmure Gordon analyst Peter Hitchens said in a conference.

Table 6: An illustration of how appeal to authority is performed. In step 1, we generate a statement using BART with the prefix “Panmure Gordon analyst Peter Hitchens confirmed that “ ”. In step 2, we move the subject and verb to the back of the sentence to diversify the sentence structure. In step 3, we swap the verb with another verb from the verb set V . In step 4, we append a preposition *in* to the sequence in step 3 and use the resulting sequence as prefix to BART’s decoder to generate the rest of the context. For step 1 and step 4, we mark the prefix sequence to the decoder in yellow, and the generated sequence in blue. To increase the diversity of the generated sequences, step 2 to step 4 are each performed 50% of the time.

weight decay to $5e-5$ and $1e-5$ for the the parameters that have been pre-trained, and $1e-3$ and $1e-3$ for other parameters. For experiments on the GROVER detector, we follow the original detection setting. GROVER is trained using Adam (Kingma and Ba, 2015) with a learning rate of $2e-5$ and a batch size of 64. Similarly, we follow the original recipe to train HDSF, which is optimized with Adam with a learning rate of $1e-2$.

All experiments are conducted on a Ubuntu 18.04 machine with NVIDIA Tesla V100. We use PyTorch 1.10.0 and Transformers 4.3.0 for constructing all models and loading pre-trained weights, except for GROVER, which operates on Tensorflow 1.13.1.

E Human Evaluation Details

In this section, we describe the survey we deliver to AMT workers for evaluating the quality of the generated articles. Annotators are presented a generated article and asked to answer a few questions regarding the quality of it. Q2 is only applicable for evaluating generated articles from PROPANEWS, in which we show the sentence that contains propaganda. The low, medium, and high ratings in the response correspond to 1, 2, and 3 scores described in Section 5.2. The questions and guidelines are illustrated as follows:

Q1: How plausible do you think of the article above?

- **Low:** It likely contains inaccurate information.
- **Medium:** Not sure.

- **High:** It unlikely contain inaccurate information.

Q2: How much does this sentence in the article affects your decision on the previous answer?

- **Low:** This sentence does not affect my answer for the previous question.
- **Medium:** This sentence somehow affect my answer for the previous question.
- **High:** This sentence largely affects my answer for the previous question.

Method	Generated Article
GROVER-GEN	Libya to restart three oilfields on Sept 1 : minister BENGHAZI , Libya , Aug 28 (Reuters) - Libya will restart two eastern oil fields in mid-September and resume crude shipment from Tobruk by the end of the month , SNCM deputy oil minister Habib Hummed said on Sunday . He said the west of Sirte was being restored and a third field in the eastern province of Brega was about to restart operations . “ Production at the Brega field has been increased by two to four percent and will resume its production on the end of the month , ” he said at a political conference . “ The second phase of production at the eastern field of Sirte came to full production by the end of the month and restarting the production from Brega field will begin in September , ” he said . He said production at the west of Sirte and east of Tobruk had already reached 95,000 barrels per day (bpd) on their main export terminals . Libya , awash with millions of dollars , most of it believed to have been siphoned off through crude theft , is pumping around 290,000 bpd in total . The first fuel terminals in Libya have reopened and are loaded with diesel and petrol and tankers from Europe are ploughing in cargo to get through a key smuggling route from Derna to the Mediterranean . Libya was sharply divided last year when armed tribesmen struck crude export terminals along its border , denying it the lucrative gas exports it was using to service southern EU countries . Libya ’s southern oil terminal was hit in September by militant attacks . (Writing by Ulf Laessing ; Editing by David Holmes)
FACTGEN	Shares Augmented and stored pipelines replaced by light pipes .Disputed concerns over biggest damage to Libya since 2003 . Despite concerns oil could be leaking . ’Russia says production is safe and it is ready to resume production . ’Greece could use more sulfur oil by 2020 . ’Energy ministers say it could be used to power two-thirds of the world economy . ’Saudi had been the biggest oil producer in the world before meltdown . ’French to continue to contain crisis . ’Russia on Libyan crisis : ‘ we are facing a serious and continuing public health crisis that is needed ’ ’But not in a vacuum of oil ’ ’World affairs . ’Greece : ‘ we need to be doing more to combat damage to the energy system . ’ ’ leftasing all oil from Libya crisis . ”
Ours (PN-SILVER)	Tue Aug 23 , 2011 4:05 pm EDT * Limited Libyan oil production could start soon * Condition of oil export terminals still unclear * Clearing pipes clogged by waxy crude could delay output By David Sheppard and Sarah Young NEW YORK/LONDON , Aug 23 -LRB- Reuters -RRB- - Limited Libyan oil production ought to resume quite quickly as most fields appear to have survived the civil war unscathed , international oil services companies say , but many are still waiting for more stability in the country before sending employees back in . In interviews with some of the biggest oil services players in pre-civil war Libya , most echoed the view that major oil fields had largely avoided serious damage during the six months of fighting , though some voiced concerns about damage to export terminals and pipelines . OPS International Chairman Gavin De Salis told Reuters Insider television that Libyan crude oil , prized for its high yield of valuable light products such as gasoline and for its low sulfur content , was quite waxy , which could clog up pipelines if they had been left unused for some time . “ There might be a little bit of effort unplugging pipelines , which is two to three months ’ worth of effort before they can resume full production , ” De Salis said . “ But that will not affect all of the pipelines or all of the fields , so they can certainly start limited production quite quickly . ” Nilsson said contacts at Libya ’s rebel oil firm Arabian Gulf Oil Company -LRB- AGOCO -RRB- informed him there had been little damage to the oilfields in the east of the country during the six-month power struggle . “ We have n’t been able to work at the oilfields during the civil war as it has not been safe , but I think within a couple of weeks we could be back to almost normal , ” Nilsson said by telephone from his office in Stockholm . “ The oil income is essential to Libya and the new government so they will want to bring it back online as soon as possible . ” Nilsson said they had several Swedish , Indian and Sudanese employees who had stayed in the country during the civil war , but total staff numbers in the country were down from around 250-300 . Nilsson said there was still a lot of work to be done in the country . De Salis said that “ a lot of damage ” had been done to Libya ’s oil infrastructure , including the destruction of some of the country ’s main oil export terminals , but he said it was too early to estimate the full extent of the damage . DAMAGE Oil firm ’s who supported the rebel government during the civil war are expected to win the lion ’s share of contracts to help relaunch the Libyan oil industry , which before the war produced some 1.6 million barrels per day of crude ...

Table 7: A qualitative comparison between generated articles from different approaches. The texts marked in orange indicate disinformation, and the texts in blue denote propaganda. We see that other approaches generate a large amount of inaccurate information, which contrasts with a property of human-written fake news mentioned in Section 1. We also note that the article generated using FACTGEN appear to be low-quality. This is likely caused by the fact that the checkpoints reported in the paper were not released and we train FACTGEN from scratch by closely following the recipe described in Shu et al. (2021). It is possible that some details of the training process of FACTGEN were missing from the paper, and hence the low generation quality.