

Political Depolarization of News Articles Using Attribute-aware Word Embeddings

Ruibo Liu,¹ Lili Wang,¹ Chenyan Jia,² Soroush Vosoughi¹

¹Department of Computer Science, Dartmouth College

²Moody College of Communication, The University of Texas at Austin
 ruibo.liu.gr@dartmouth.edu, soroush.vosoughi@dartmouth.edu

Abstract

Political polarization in the US is on the rise. This polarization negatively affects the public sphere by contributing to the creation of ideological echo chambers. In this paper, we focus on addressing one of the factors that contributes to this polarity, polarized media. We introduce a framework for depolarizing news articles. Given an article on a certain topic with a particular ideological slant (e.g., liberal or conservative), the framework first detects polar language in the article and then generates a new article with the polar language replaced with neutral expressions. To detect polar words, we train a multi-attribute-aware word embedding model that is aware of ideology and topics on 360k full-length media articles. Then, for text generation, we propose a new algorithm called Text Annealing Depolarization Algorithm (TADA). TADA retrieves neutral expressions from the word embedding model that not only decrease ideological polarity but also preserve the original argument of the text, while maintaining grammatical correctness. We evaluate our framework by comparing the depolarized output of our model in two modes, fully-automatic and semi-automatic, on 99 stories spanning 11 topics. Based on feedback from 161 human testers, our framework successfully depolarized 90.1% of paragraphs in semi-automatic mode and 78.3% of paragraphs in fully-automatic mode. Furthermore, 81.2% of the testers agree that the non-polar content information is well-preserved and 79% agree that depolarization does not harm semantic correctness when they compare the original text and the depolarized text. Our work shows that data-driven methods can help to locate political polarity and aid in the depolarization of articles.

Introduction

Political polarization refers to an individual’s stance on a given issue that is heavily affected by their identification with a particular political party (e.g., Democratic or Republican) or ideology (e.g., liberal or conservative). Political polarization has been recognized as a major issue undermining western democracies (Youngmann 2019; Groseclose and Milyo 2005; Gentzkow and Shapiro 2010). Partisan media often exacerbates this problem; in fact, it is reported that politically polarized articles can gradually change the attitudes of communities (Iyyer et al. 2014), limit the topics in public discussions (Chen et al. 2018b), or even have a substantial effect on the evolution of policy (Dardis et al. 2008).

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this paper, we propose a data-driven framework to detect and reduce political polarity in a given news article. The framework is intended to be used by authors to depolarize their articles while maintaining their message. The tool will not change the stance of an author; it will transform the text to say the same thing, but in a less polar way. To illustrate this point, consider an article talking about illegal immigrants. The polarity of that article can be changed by simply changing the phrase *illegal immigrants* with *undocumented immigrants*. The polarity can be changed the other way around by changing the former phrase with *illegal aliens*. This change in polarity does not affect the message of the article; it just softens (in the case of the first example) or hardens (in the case of the second example) the tone. The polarity of an article can be transferred by changing a few polarity markers without changing the meaning of the text; these polarity markers can be explicit phrases (e.g., *illegal immigrants*) or more subtle expressions (e.g., *can’t pay for health insurance* vs. *can’t prove coverage of health insurance*).

Our framework is composed of two main parts: a polarity detection mechanism based on attribute-aware word embeddings, and an algorithm for fully automatic or semi-automatic depolarization. The polarity detection model detects phrases and expressions that make an article polar. The model is context-aware (topic-aware), meaning that a phrase could be detected as polar in articles about a particular topic while not being detected as polar in other topics. The depolarization algorithm suggests changes to the polar language detected by the polarity detection model that would reduce the polarity of the article, while at the same time preserving the meaning, the message and the stance of the article (e.g., the algorithm would not change the word *ban* to *approve* as that would change the stance of the article). Furthermore, the depolarization algorithm attempts to keep the article contextually meaningful, cohesive, comprehensible and grammatically correct (something that most current transfer models do not do well).

Our method is inspired by recent works on encoding attribute-specific information into word embeddings (Subramanian et al. 2018; Lample et al. 2019; Logeswaran, Lee, and Bengio 2018), and works on adversarial text generation

(Ren et al. 2019; Alzantot et al. 2018; Chen et al. 2018a; Zhao et al. 2018; Chen et al. 2017; Singh and Palod 2018). In our case, we treat political ideology and topic as two separate attributes and train a multi-attribute-aware word embedding model. Our embeddings can not only capture the general context of words, but also the nuanced differences in their usage in specific contexts, such as different topics and ideological slants. Next, we propose an algorithm, called TADA, for generating depolarized (or neutral) text that can replace polar language (identified using the multi-attribute-aware word embedding model) with grammatically correct *neutral* language (also extracted using the our multi-attribute-aware model). TADA can be used fully automatically, where a new depolarized article is fully generated by the tool, or semi-automatically, where users can decide which of the suggested changes to accept. We evaluate our models and algorithm qualitatively and quantitatively.

Related Work

Political polarity and bias in the media have long been studied in the fields of communication and political science, going back to at least the 1950s (Scammell and Semetko 2018). Its impact on the political process has also been extensively studied (Kahneman and Tversky 2013; Huillery 2009; Glaeser and Goldin 2007). Specifically, studies show that if news consumers regularly receive polar news, they are prone to adopting similar polar views, and tend to start exclusively “following” the media outlets whose reporting conforms with their established beliefs, thus spiralling into an echo chamber where their bias is further reinforced.

Polarity Detection There is a wealth of work on how to detect polarity in media. Yano et al. (Yano, Resnik, and Smith 2010) performed sentence-level manual polarity annotation on blog posts from 2008 and they showed that words capturing emotion and named entities of opposing political parties can be a strong indicator of polarity. In recent years, automated identification of political polarity and bias in news has gained more attention. Gentzkow et al. (Gentzkow and Shapiro 2010) derive a “slant index” to rate the ideological polarity of newspapers. They first identify words that are used much more frequently by one party than by another from the Congressional Record, and then compute a slant index (same concept as polarity) based on the frequency of these party-specific tokens. Lin et al. (Lin, Bagrow, and Lazer 2011) proposed a scheme for bias categorization. The scheme includes the political party, frequently mentioned legislators, region, ideology, and gender. Iyyer et al. (Iyyer et al. 2014) apply a recursive neural network (RNN) to identify the political position evinced by a sentence. Owing to RNN’s modelling power on sequential data, they claim that their model can capture both syntactic and semantic features. However, in terms of the detection output of the highest probability n-grams for two ideologies, their model struggles to distinguish non-polar content and polar indicators. Lexical variations have also been studied in non-political contexts. Shoemark et al. (Shoemark et al. 2017) conduct a large-scale study of dialect variation on Twitter in the UK. Their data-driven approach identifies Scotland-specific lexi-

cal features.

Text Style Transfer There are parallels between the depolarization stage of our framework and the general field of linguistic style transfer. One of the mainstream methods for style transfer is to rely on variants of auto-encoders to generate “transferred” text (Zhou and Neubig 2017; Mueller, Gifford, and Jaakkola 2017; Guu et al. 2018; Yang and Eisenstein 2017; Bowman et al. 2015). Works that rely on auto-encoders for style transfer require a large parallel corpus. There is a lack of ideal parallel corpus for our task that aligns sentences with the same content but different political ideology (Subramanian et al. 2018; Tian, Hu, and Yu 2018). Due to the lack of such corpus and also because of the difficulty of disentangling polar text with non-polar content (Dai et al. 2019a; Prabhumoye et al. 2018), these works fail to output meaningful indicators of polar text and in most cases, the output is problematic in grammar and hardly understandable (Wu et al. 2019). Chen et al.’s (Chen et al. 2018b) work, which uses style transfer to flip the bias in the news headlines, is the most similar to ours. They first analyze bias among the articles and then use a cross-aligned auto-encoder trained on opposite-ideology news data to generate flipped titles. The main limitation of this work, as stated by the authors, is that even in their successful cases, the overlap of generated and ground-truth headlines is very low, which means the auto-encoder model has the tendency to discard too much information even from non-polar content part. We differ from their work in the model selection and also integrating salient words up-sampling into the feature selection (this is explained in detail later in the paper). Also, we extend the task to the depolarization of the entire article, instead of only headlines, which is a more challenging problem since we have to preserve contextual and grammatical correctness across the whole article.

More recent works on style transfer pay more attention to the quality of generation and whether the newly created text can successfully preserve the non-style content. These approaches are mostly template-based: they first identify target words that contribute a lot to style, and then replace these words and generate new expressions. Li et al. (Li et al. 2018) propose a simple but efficient approach called “Delete, Retrieve, Generate”, which deletes stylized n-grams based on corpus-level statistics and stylizes it based on similar, retrieved sentences. Wu et al.’s work “Point-Then-Operate (PTO)” (Wu et al. 2019) also follows this two-step manner to transfer style, which consists of a high-level agent that proposes operation positions and a low-level agent that alters the sentence. Huang et al. (Huang et al. 2019) bring synonyms from a dictionary as external reference information to improve generation quality, and they view the transfer generation task from the point of view of human’s: humans tend to replace words or phrases in the original sentence with their corresponding synonyms and make necessary changes to ensure the new sentences are fluent and grammatically correct. Our work is inspired and builds upon these dictionary-guided methods, but rather than using an off-the-shelf dictionary that can only guarantee grammatical correctness, we train a “dynamic dictionary” (dynamic with

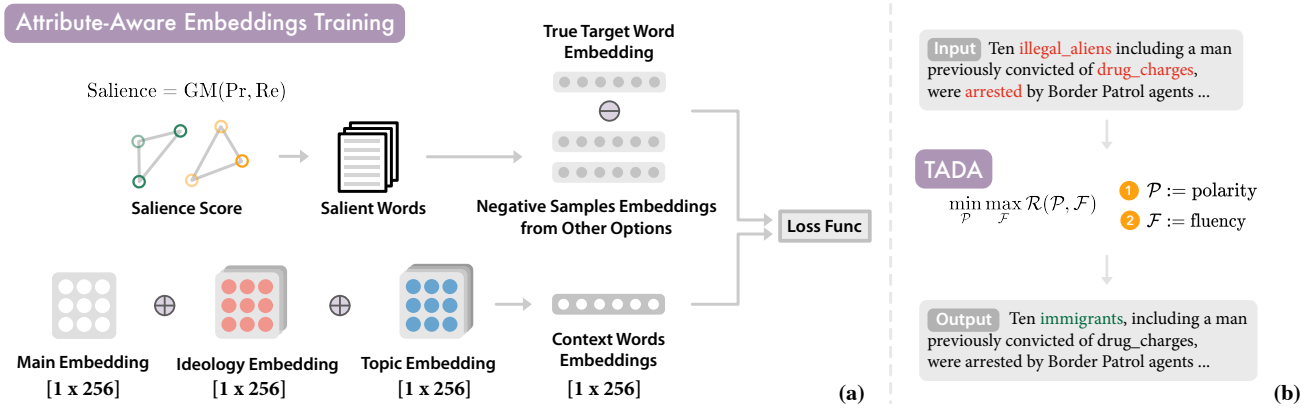


Figure 1: Overview of (a) dynamic attribute-aware embeddings training, and (b) TADA algorithm of our framework.

respect to topic and ideology) on an attributes-labeled corpus so that such a dictionary can not only preserve grammar but also can serve as a strong reference for a depolarized-text generation. This paper is related to another attempt from our lab at transferring the political polarity of news articles using a transformer-based conditional generator (Liu, Jia, and Vosoughi 2021).

Approach

Our model is comprised of a dynamic multi-attribute embedding model which is used to detect polar phrases and a probabilistic algorithm for depolarizing the detected polar text. In this section we explain our framework in detail.

Dynamic Attribute-aware Model

In order to capture attribute-invariant as well as multi-attribute-aware information, we extend a current single-attribute dynamic model (Gillani and Levy 2019) to multi-attribute case through making following changes: (1) we adopt a two-step procedure, first we pre-train attribute-specific embeddings using datasets for each attribute (i.e., we train embedding models using articles with different ideologies and topics). Next, we jointly train an attribute-invariant embedding on the full dataset with attribute weights initialised as pre-trained embeddings, (2) we integrate salient words (words that can help distinguish one attribute option from another) features into the joint training step to promote saliency of attribute-specific words, (3) we replace the original negative sampling objective function with attribute-aware negative sampling objective so that our training can converge faster and the model is able to reflect attribute features in the final embeddings.

The basic idea of our attribute-aware model for the embedding of a word w can be defined as follows:

$$E(w, A_1, \dots, A_k) = E_m(w) + \sum_{A_1} p_{A_1} E_{A_1}(w) + \dots + \sum_{A_k} p_{A_k} E_{A_k}(w) \quad (1)$$

where $E_m \in \mathbb{R}^{|V| \times \text{emb}}$ denotes the attribute-invariant embedding (main embedding) of w , and $E_{A_1}(w), \dots, E_{A_k}(w) \in \mathbb{R}^{|V| \times k}$ represents the shift of main embedding in specific attribute domain. We denote the embedding corresponding to each possible option for an attribute A as $E_A(w)$; note that this framework supports embeddings for multiple attributes (denoted by $E_{A_1}(w)$ to $E_{A_k}(w)$). For each option in each attribute, we sum up all embeddings from all options in a weighted form (p_{A_i} is the weight vector for attribute i), and add it to the main embedding. We add these embeddings along the embedding axis and use the new summed-up matrix ($E(w, A_1, \dots, A_k)$) as the final embedding. The resulting model is shown in Figure 1 (a).

Integrating Salient Words One limitation of our “main + attribute” structure is that all negative samples that help training converge come from the same vocabulary set, meaning that we cannot have the model pay more attention to words that are more salient in one category (within the same attribute). This becomes an issue when an attribute has many categories (for instance our topic attribute which has 11 categories). As shown in Table 1, the most similar terms of the same query word *illegal* show nearly no difference when we switch the topic.

Certain words are strongly associated with a particular attribute. For instance the phrase *illegal aliens* is a strong indicator of conservative ideology, which means that it should be a good candidate of negative samples in the other two ideology options. Based on this observation, we identify a set of salient words that serve as strong indicators of different attribute options. We integrate these salient words into our training by making them more likely to appear in negative samples. This way our attribute-aware model can converge faster and output higher-quality word embeddings. Specifically, the saliency score \mathcal{S} for a word w with respect to the attribute option O_i is:

$$S(w, O_i) = \text{GM}\left(\frac{\text{count}(w, O_i)}{\sum_{n=1}^N \text{count}(w, O_n)}, \frac{\text{count}(w, O_i)}{\sum_{w_i \in |V|} \text{count}(w_i, O_i)}\right) \quad (2)$$

where the first fraction is the probability that a word w is from option O_i (i.e., $\mathbb{P}(O_i|w)$), while the second fraction is the probability that w can be found in O_i (i.e., $\mathbb{P}(w|O_i)$). GM is the geometry mean of these two terms. Note that we perform laplace smoothing and normal distribution normalization on the two terms.

Using this definition, the salience score of a word w with respect to option O_i depends on $\mathbb{P}(w|O_i)$ and $\mathbb{P}(O_i|w)$. This means that a word with a high salience score is highly frequent in O_i while not being universally frequent (similar to the TF-IDF score). We use this score to sort our vocabulary set for each option, and then we apply a weighted (in terms of salient score) up-sampling on those words that are marked as salient (we explain this in more detail in the following section). Specifically, when we generate negative samples for a given activated option of an attribute (e.g., the option liberal for ideology attribute), we randomly pick words with a certain probability that has been modified by the salient words up-sampling algorithm.

Attribute-aware Objective Function We choose a CBOW (Continuous Bag of Words) model (Mikolov et al. 2013) as our base to train the attribute embedding since it is reported to perform relatively well with large datasets (Schnabel et al. 2015). The original negative sampling objective function from Mikolov’s work (Mikolov et al. 2013) can be represented as:

$$J(\theta) = \arg \max_{\theta} (\log \sigma(v_c \cdot v_w) + \sum_{(w,c) \in D'} \log \sigma(-v_c \cdot v_w)) \quad (3)$$

where v_w and v_c are a pair of word and context in vector form, σ denotes the sigmoid function and D' is the set of so-called “negative samples”. The goal is to find parameters to maximize the probability of positive cases (context words c are indeed around word w) while minimize the probability of negative cases (word-context pairs $(w, c) \in D'$ where D' are a set of randomly sampled pairs). This works well if no attribute is considered. However, it is highly possible that some negative cases we randomly sample for one attribute are actually positive samples in another attribute. As described earlier, this problem becomes even more severe as the number of attributes and the options for the attributes increases.

In the previous section we described how we create a list of salient words for each attribute option. We integrate these words in the objective function to improve negative sampling and dampen the issues described above. We do this by creating a “reversed” salient-word bank for each option O_i , which is in turn used to generate negative samples. Words that are salient with respect to an option appear less often in the reversed bank of that option and vice versa. The count

for word w in the reversed salient-word bank for option O_i is calculated as:

$$\text{Count}(w) = M \times \frac{1}{p}, \text{ where } p = \frac{S(w, O_i)}{\sum_{w \in V} S(w, O_i)} \quad (4)$$

here $S(w, O_i)$ is the salience score we defined in Equation 2, M is a relatively large number (10^8) close to the total count of words in our dataset. This equation ensures that words that are salient with respect to an attribute option will appear less in that options word bank and thus has less chance of being selected as a negative sample and vice versa. We generate and store these words banks in advance to speed training.

Through this salient-based negative sampling method, we are able to improve the quality of our attribute-aware word embeddings. Our model is evaluated later in this paper.

Detect and Depolarize Polar Text

Detect Polar Words Using the attribute-aware model, we define a measure of polarity for a word as the distance between its embeddings in different ideology attributes (liberal, neutral, conservative). Mathematically, our definition of polarity score \mathcal{P} for each word w in topic t is:

$$\mathcal{P}(w, t) = \sum_{s_i, s_j \in S} \cos(E(w, s_i, t), E(w, s_j, t)) \quad (5)$$

where S is the set of all ideology attributes and E is the attribute-aware embedding. We apply z-score normalization on each word’s polarity score within the same topic. As a result all polarity scores’ mean μ is located at zero, with several sigma σ deviation in two directions. We treat the words that have above-zero polarity as polar words, and for the sake of generalization we filter out words whose frequency is lower than 500.

The polarity of a paragraph (or even a full article) can be calculated by adding the polarity score of each polar word (a word with normalized polarity greater than 0). We ignore non-polar words in our calculation as we do not want the polarity of a paragraph to be artificially lowered due to its length. At this point, we have the polarity score of each paragraph and a set of polar words in the paragraphs. The next step is to replace those polar words with neutral words such that the polarity of the paragraphs go down.

Depolarized-Text Generation In this section, we introduce the Text Annealing Depolarization Algorithm (TADA) as our solution to the depolarization task. The workflow of TADA is shown in Figure 1 (b).

In addition to polarity detection, the multi-attribute-aware word embedding model can also be used to find a polar word’s neutral replacement.

For a word w_i from source ideology s_i within topic t , its neutral expression $w_{\mathcal{N}}$ can be retrieved using the following equation:

$$w_{\mathcal{N}} := \arg \min(\cos(E(w_i, s_i, t), E(w_i, \text{neutral}, t))) \quad (6)$$

this amounts to retrieving w_i 's most similar word in the neutral ideology and the same topic t . The distance is measured using cosine similarity. To ensure the correctness of grammar, we filter out those candidates whose part-of-speech (POS) are not the same as the original word's POS. For each detected polar word, a neutral candidate list (denoted as $w_{\mathcal{N}}$) is retrieved and sorted using this method. We limit the list for each polar word to the top 20 candidates for the search algorithm (which means $|w_{\mathcal{N}}| = 20$).

TADA is a modified simulated annealing algorithm that balances depolarization performance and fluency. The algorithm is also highly efficient. Let us first define the optimization goal of our task: Given a paragraph consisting of words \bar{w} in topic t , assume after s replacement steps we obtain a new sequence of words $\bar{w}^* = \{w_1, \dots, w_m\}$. The reward \mathcal{R} of the whole modification is:

$$\mathcal{R}(\bar{w}^*, \bar{w}, t) = \frac{1}{s + \lambda} \left[\underbrace{-\mathcal{P}(\bar{w}^*, t)}_{\text{step 1}} + \underbrace{\mathcal{F}(\bar{w}^*, \bar{w})}_{\text{step 2}} \right] \quad (7)$$

where \mathcal{P} is the polarity score we mentioned in the previous part, \mathcal{F} is the fluency score of the newly generated sequence \bar{w}^* with respect to the original paragraph \bar{w} . Also, we use s to normalize the total rewards in terms of how many steps it takes to reach the optimum replacement, and we add a regularization term λ to avoid infinity reward \mathcal{R} caused by zero step/no modification (i.e., $s = 0$). We set the hyperparameter $\lambda = 0.01$ through empirical observation.

As shown in Equation 7, we compute the reward in a two-step manner: in step one, we compute the polarity score \mathcal{P} of the randomly generated modified sequences (\bar{w}^*) (note that a modification is only made when the polarity score of the sentence is lowered); and then in step two, we sum up the fluency score \mathcal{F} and the negative polarity score of each candidate sequence to produce the final reward \mathcal{R} . Given such a reward definition, the training goal of our TADA algorithm is to pick the proper sequence modification that has the highest reward after several iterations.

It is possible that no modifications will be made in the first step if there are none that reduce the polarity score. In that case the second step will not go into effect and the algorithm will terminate with no changes on the sequence. To decrease the number of unchanged cases we can set a lower T_{\min} in Algorithm 1, which leads to more iterations for the depolarization.

We choose the BLEU (Papineni et al. 2002) score as an estimation of fluency (\mathcal{F}). BLEU is computed based on how many n-gram overlap between the newly generated text and the original one. We believe it is more suitable than other language model scoring methods because the replacement (i.e., depolarization) causes a slight stylistic shift (words with the same meaning but different styles). Compared to other language model methods (e.g., BERT) BLEU is more sensitive to such changes. Another reason for selecting BLEU is speed. Some neural language models, like BERT, suffer from slow evaluation scoring (Lan et al. 2019). Since we are not rewriting the whole sentence but perform-

ing modification on the tokens, we choose BLEU as our light-weight and fast fluency metric.

Algorithm 1: Text Annealing Depolarization

Input: Original polar text $\langle w_1, \dots, w_m \rangle$
Output: Neutral text *neutral*
 $T \leftarrow T_{\max}$
 $neutral \leftarrow []$
for w_1 **to** w_m **do**
 if $\mathcal{P}(w_i) > \text{threshold}(\mathcal{P}(w_i)$ by Eq. 5) **then**
 retrieve $w_{\mathcal{N}}$ of w_i (by Eq. 6)
 $t \leftarrow 0$
while $T > T_{\min}$ **do**
 Randomly init *best* from $w_{\mathcal{N}}$
 $next \leftarrow \text{RandomPickCandidates}(T, best)$
 $\Delta \mathcal{R} \leftarrow \mathcal{R}(next) - \mathcal{R}(best)$ (\mathcal{R} by Eq. 7)
 if $\Delta \mathcal{R} < 0 \wedge \text{random} < \mathbb{P}_{\text{accept}}(T, \Delta \mathcal{R})$ **then**
 $best \leftarrow next$
 else
 Discard *next*
 $t \leftarrow t + 1$
 $T \leftarrow \frac{T_{\max}}{\log(t+d)}$
 $neutral \leftarrow \text{Merge } best \text{ back to text}$
return *neutral*

Given the reward function, the fully-automatic mode relies on an efficient search algorithm that can find the best combination of neutral words. Given an article whose length is l and the vocab size is $|V|$ (a very large number), the search space of common sequential decoding methods (Shen et al. 2017; John et al. 2018; Dai et al. 2019b) is $|V|^l$, and the final generation often drops into a local optimum because of its accumulation feature. Reinforcement-learning-based methods (Gong et al. 2019; Xu et al. 2018) have more accurate quality control over the generation and partially overcome the local optimum problem, however, these methods rely on continuous sampling of candidate sequences (still $|V|^l$ search space), which result in low efficiency. Our algorithm, TADA, improves the efficiency in two perspectives: We first narrow the search space for each token from $|V|$ to $|w_{\mathcal{N}}|$, which is a limited-length neutral replacement words list proposed by our dynamic attribute-aware model. Then, instead of using sequential decoding to generate a replaced sequence, we use the simulated annealing algorithm as our template, and modify it to meet our requirements. The overview of our algorithm (called TADA) is shown in Algorithm 1.

We keep the general framework of the classical simulated annealing algorithm: The optimization procedure begins with a maximum temperature T_{\max} , and after several iterations the optimization stops, when the temperature reaches T_{\min} . Instead of relying on the normal token-by-token decision, we decouple the search space with the length of the input. The replacement only happens when there are polar words detected, and we deploy a propose-and-judge manner as the core logic of TADA. In this way, assuming

there are k words being detected as polar words, the search space of our algorithm is thus $|w_{\mathcal{N}}|^k$, where k is not necessarily related to the input length and $|w_{\mathcal{N}}|$ is limited to 20. We make use of the efficient search method in the classical simulated annealing algorithm. Each iteration of our algorithm can be described in three steps:

1. Randomly pick polar words in a paragraph, and randomly pick neutral candidates for each selected polar word. Use these candidates to do the replacement, and compute the reward score (\mathcal{R}).
2. If the newly computed reward score is smaller than the best, which means the polarity has decreased and fluency is well-conserved, accept these candidates to generate a new paragraph.
3. If not, accept the candidates with probability $\mathbb{P}_{\text{accept}} = \exp(-\frac{\Delta\mathcal{R}}{T})$.

We also use several tricks to improve the generation quality. For example, in order to guarantee that the picked candidates are grammatically correct, we filter out the candidates whose POS is not the same as the POS of the original word. Also, we dynamically adjust the “cooling speed” (ΔT) based on the length of the paragraph since longer paragraph may need more iterations to find the best candidates combination. In general, given the observation that certain polar words or phrases contribute a lot to the polarity, TADA adopts a propose-and-judge procedure, avoiding the unnecessary token-by-token sequential decoding from scratch. We make use of a pre-trained attribute-aware model to shrink the search space of neutral replacement to guarantee our algorithm runs efficiently. (All the parameters used in the model are listed in the Implementation section.)

Implementation

Datasets

The data for this project was collected from Media Cloud¹, which is an academic research project led by MIT and Harvard University. The platform collects articles from a large number of media outlets and makes the content available via an API.

We followed the political issues defined by a survey-based website² and chose 11 topics (shown in Table 4) for our data collection. For each topic, we use the words under the corresponding title on isidewith.com as query keywords. (E.g. topic: Social Issue \rightarrow query terms: *abortion, gay marriage, death penalty*, etc) We collected around 360k full-length articles from May 1st, 2018 to May 1st, 2019 (about 6GB of plain text), as our training data (similar to Liu et al. 2021). The articles are from 22 media outlets. We assign an ideological polarity label to each outlet using data from the Pew Research Center. The data from Pew is based on a survey of news consumption by people with different political affiliation³. The labels we assigned for polarity are *liberal*

¹<https://mediacloud.org/>

²<https://www.isidewith.com/polls>

³<https://www.journalism.org/interactives/media-polarization/table/overall/>

(corresponding to Pew’s *consistently liberal* and *mostly liberal* labels), *neutral* (corresponding to Pew’s *mixed* label), and *conservative* (corresponding to Pew’s *consistently conservative* and *mostly conservative* labels). Each article was also labeled with a topic. The assigned topic labels are based on the corresponding title name for the query terms. For example an article which was collected by the *abortion* keyword query would be labeled as having the topic Social Issues.

Training

We train our attribute-aware model in two steps: (1) We train the attribute embeddings with the part of the training data that correspond to the attribute option that is activated. For example, the data we use to train liberal embeddings is all the articles from liberal media (including New York Times, NPR, Huffington Post, etc.), (2) After we obtain the embeddings for all possible attribute options (in our case, there are three ideology embeddings and 11 topic embeddings, totaling 14 embeddings), we jointly train the main (universal) embedding and all the attribute embeddings. This step is crucial as it aligns all the words from different semantic spaces into the same vector space.

The training of our dynamic attribute-aware model on a Google Cloud Platform with NVIDIA Tesla P100 GPU took about one hour for five epochs on the joint training (after five epochs we observed a decrease in the accuracy, so we choose five as the epoch number). The training accuracy reached around 87% at the 5th epoch. All our embeddings have 256 dimensions. The number of negative samples is 32. We add L2 regularizer 10^{-8} to the fused embedding layer (main + attribute embeddings). The window size of CBOW is 3.

For the TADA algorithm, we set the start temperature T_{max} to 1000 and set T_{min} to 100. We use t to record the step sizes, and at the end of each iteration T will be updated to $\frac{T_{\text{max}}}{\log(t+d)}$ where d can be set as static constant or a length-related variable. We set d to 1 for the efficiency requirement.

Two Modes: Semi and Fully-automatic

There are two core modules in our framework, the dynamic word embedding model and TADA depolarization algorithm. For the purpose of ablation study, we develop two modes of our framework: semi-automatic mode and fully automatic mode. Fully automatic mode makes use of the whole framework, while the semi-automatic mode only uses a partial TADA algorithm requiring the help of human editors to depolarize the article. Specifically, after our system detects the polar words in a given article, it retrieves their corresponding neutral options and provide them for human editors to choose from. The human editors have to then consider various criteria, such as the contextual consistency, the overall fluency, and the depolarization strength when performing the replacement. Ideally, we expect the TADA algorithm to produce similar results as the human-guided version.

Evaluation

At its core, our framework is comprised of two main parts: the attribute-aware word embeddings, and the TADA depo-




Figure 2: Several visualization examples in embedding space to verify our model manages to capture attribute-aware polarity. In (a), we plot 2D projections of the words “government” and “washington” from liberal, neutral and conservative within the same topic “Immigration”. In (b) and (c), we plot the ten nearest words to the word “illegal” in two different topics (“Immigration” and “Criminal Justice”). We prove that our model can not only measure the lexical polarity within the same topic (ideology-aware), but also reflect the discrepancy between different topics (topic-aware).

larization algorithm. Thus, our evaluation is focused on (1) Illustrating that our word embeddings manage to capture attribute information, (2) Measuring the effectiveness of our depolarization algorithm, TADA.

Topic: Immigration		Topic: Healthcare	
Term: illegal			
Word	Cosine	Word	Cosine
<i>illegally</i>	0.2714	<i>illegally</i>	0.2203
<i>possession</i>	0.2163	<i>banning</i>	0.1998
<i>criminal</i>	0.2056	<i>possession</i>	0.1942
<i>illicit</i>	0.2016	<i>smuggle</i> ◁	0.1906
<i>suspected</i>	0.2013	<i>border_crossing</i> ★	0.1888
<i>unlawful</i>	0.1996	<i>illicit</i>	0.1878
<i>prohibited</i>	0.1942	<i>prohibited</i>	0.1854
<i>drugs</i> ★	0.1935	<i>trafficking</i> ◁	0.1845
<i>felony</i>	0.1907	<i>drugs</i> ★	0.1754

Table 1: The top 10 most similar words to the word *illegal* in the immigration and drugs topics, *before* we integrate salient words into the joint training. We use ◁ to mark words that are unique to one topic and we use ★ to mark the phrase *border_crossing* and the word *drug* which we think should not appear in the most similar list for both topics.

Attribute-Aware Word Embeddings

To illustrate that our word embeddings successfully learned ideology and topic information, we first perform a qualitative experiment to test our model’s awareness of attributes. We choose a pair of words *washington* vs. *government*, which has been reported as an example of how the left and the right use different words to refer to the same thing (Abadi 2017; Miller 2011), to check whether our attribute-aware embeddings can successfully discover such known knowledge. Since the second step of our training involves jointly training attribute-invariant and attribute-aware embeddings, we are able to extract (and plot) word vectors from different attribute domains in the same vector space. Figure 2 shows the 2D projections of the words *washington* and *government*

from different ideological attributes (but the same topic, in this case, “Immigration”).

Topic: Immigration		Topic: Healthcare	
Term: illegal			
Word	Cosine	Word	Cosine
<i>illegally</i>	0.2872	<i>illegally</i>	0.2682
<i>unlawful</i>	0.2422	<i>possession</i>	0.2266
<i>trafficking</i>	0.2311	<i>criminal</i>	0.1988
<i>prohibited</i>	0.2217	<i>drugs</i> ↑	0.1979
<i>banning</i>	0.2208	<i>money_laundering</i> ↑	0.1966
<i>minors</i>	0.2101	<i>smuggle</i>	0.1917
<i>felony</i>	0.2099	<i>unlawful</i>	0.1904
<i>undocumented</i> ↑	0.2081	<i>marijuana</i> ↑	0.1874
<i>firearms</i>	0.2056	<i>suspected</i>	0.1848

Table 2: The top 10 most similar words to the word *illegal* in the “Immigration” and “Drugs” topics, *after* we integrate salient words into the joint training. We use ↑ to indicate the rise in ranking of certain words.

Recall we defined the polarity of a word as the distance between its embeddings in different ideology attributes. This corresponds to the area of the triangle shown in figure 2. We can see that the word *washington* is more polar than the word *government*. This makes sense intuitively as the word *washington* is a political term often used to refer to the federal government, like phrases *Washington spending*, *Washington waste*, *Washington taxation*, which is reported in the book *Words that Work* by Frank Luntz (Luntz 2007) and other related works (Marangos 2008; Lemmon 2013). In this qualitative case we thus verify our model is able to correctly capture the polarity phenomenon.

In the second experiment, we look at the similarity task for a given word in different topics. As shown in Table 1 and Table 2, we perform similarity checks on the term *illegal* within two topics: “Immigration”, and “Healthcare”, and we use cosine similarity as our distance metric. Table 1 shows the most similar words to the word *illegal* before we integrate salient words into the training. Though there are some differences between the two topics, it is hard to

Topics	Liberal (L)			Neutral (N)		Conservative (C)			Total Polar		Success %
	before	after	$\Delta\% \downarrow$	before	after	before	after	$\Delta\% \downarrow$	before	after	
Social Issues	7	3	57.1%	4	12	7	3	57.1%	14	6	57.1%
Immigration	12	4	60.1%	11	29	20	10	50.0%	32	14	56.3%
Miscellaneous	33	12	63.3%	10	74	60	17	71.7%	93	29	67.0%
Healthcare	20	2	90.0%	9	47	33	13	60.0%	53	15	71.7%
Climate	25	13	48.0%	4	33	30	13	56.7%	55	26	52.8%
Trade	79	24	69.6%	40	165	91	21	76.9%	170	45	73.5%
Election	106	34	67.9%	40	207	142	47	64.6%	248	81	67.3%
Foreign Policy	28	15	46.4%	11	70	67	21	68.7%	95	36	62.1%
Education	39	17	56.4%	15	51	32	18	43.8%	71	35	50.7%
Criminal Justice	39	4	89.7%	14	71	36	14	61.0%	75	18	76.0%
Vaccination	20	3	85.0%	5	30	11	3	72.2%	31	6	80.6%
Overall	408	131	67.9%	163	789	529	180	66.0%	937	311	66.8%

Table 3: The success rate of fooling an external ideology classifier after we perform the depolarization algorithm TADA on 1,100 original polar text across 11 topics. The success rate is defined by Equation 8. The F1 score of the external classifier is 0.8974 (liberal, neutral, conservative three-class classification).

tell the difference between these two topics. Also the word *border_crossing* is sort of misleading in the context of drugs: Though some drug trade is related to crossing the border, the phrase is more likely to be used in the context of immigration. At the very least, if such a phrase is frequent in both topics indeed, we would like to see the model treat them semantically different in different topics. Another potentially problematic word is *drug*, which ranks high in both of the topics. We compare these results to ones generated using a model that uses our salient word up-sampling algorithm.

The version of our model that uses salient word up-sampling performs better on the same similarity task. Table 2 shows the results. We can observe that words like *drugs*, *money_laundering* and *marijuana* rank higher in the list for the topic drug, and the word *undocumented* has joined the list for the topic immigration. Though qualitative in nature, these evaluations help shed light on the performance of different aspects of our model.

Depolarized-Text Generation Algorithm (TADA)

To verify the effectiveness of our depolarization algorithm (TADA), we use an automatic evaluation framework. We train an ideology classifier with three classes, liberal, neutral, and conservative on the whole training dataset ($F1 \approx 90\%$) to be used for judgement. For a given text, we use this classifier to detect its ideology before and after it has been depolarized by TADA. For a given corpus of text (comprised of liberal (L), neutral (N) and conservative (C) texts), we define the depolarization success rate of TADA as:

$$\text{Success \%} = \frac{\text{count}(\text{after N}) - \text{count}(\text{before N})}{\text{count}(\text{before L}) + \text{count}(\text{before C})} \quad (8)$$

The rate focuses on the increased number of neutral text against the total number of polar text by measuring how

many newly generated samples will be judged by the external classifier to be neutral. We choose fasttext by Facebook (Bojanowski et al. 2016) as our external classifier and pick the best parameter set through grid search (dimension=500, n-gram=3). For our evaluation, we sample 1,100 polar articles (L and C) from 11 topics, matching the original distribution. We run these article through TADA and measure the success rate using Equation 8.

The result is shown in Table 3. We list all 11 topics results and show the before and after ideology distribution. TADA performs slightly better on the originally liberal text (overall 67.9% vs. 66.0% of originally conservative), and from the topic-wise perspective, TADA performs best on the Vaccination topic (80.6%) but worst on the Education topic (50.7%). We think that the discrepancy between topics is due to certain topics having more salient words and phrases, which gives TADA more options to work with, and some topics having very few salient phrases, which causes potential difficulties for the TADA algorithm to depolarize. Overall, we achieve a success rate of 66.8% over the 1,100 samples.

Human Judgement Evaluation

To figure out whether TADA’s depolarization is truly recognized as successful by humans (which is the ultimate goal of this paper), we conducted a online experiment on the Amazon Mechanical Turk (MTurk) platform in 2019. MTurk is an increasingly popular platform for online data collection with a population of workers that is typically more diverse in age and racial distribution than American college samples (Buhrmester, Kwang, and Gosling 2011).

Participants

We recruited 161 participants ($N = 161$) to take part in the evaluation process. The participants were all from the United

States. The average age of the participants was 40.55 years-old ($SD=13.99$, $Median=37$). More than half (57.1%) of the participants were male, and 42.9% were female. Participants on average received 15.8 years of education ($SD = 2.80$, $Median = 16$). When asked to self-report their party affiliations, 83 participants (51.5%) were self-reported Democrats, 48 participants (29.8%) were self-reported Republican, and 30 participants (18.6%) were independent. Accordingly, 78 participants (48.4%) preferred to read liberal media, and 42 participants (26.1%) preferred to read conservative media.

Stimuli

We randomly selected 33 sets of paragraphs from TADA's generation pool from 11 topics (listed in Table 4). Each set has one original paragraph, one paragraph with text edited by semi-automatic mode and one with text edited by the fully-automatic mode. In total, we selected 99 paragraphs as our evaluation samples.

Procedures

The 161 participants were randomly assigned to one of the three groups. People who were assigned to group A were presented texts from the topics Social Issues, Immigration, Miscellaneous, Healthcare. Group B was shown the topics Climate, Trade, Elections, and Foreign Policy and group C the topics Education, Criminal Justice, Vaccination. We did this to ensure that single participant does not have to read texts pertaining to 11 different topics.

Before the experiment, participants were asked to answer a pre-exposure questionnaire about their demographics, party affiliation, political attitude, and favorite media outlets. During the experiment, for each set, participants were exposed to three versions of paragraphs (original, depolarized using the semi-automatic mode, and depolarized using the fully-automatic mode) without being informed of the actual type. After reading each paragraph without knowing whether it was depolarized or not, participants were asked to rate the polarity, ideology and readability of the paragraphs on the 7-point Likert scale.

Next, the participants were informed of the actual type of each version and were presented the three versions of the stories at the same time. The participants were then asked to answer a post-exposure questionnaire about the performance of the algorithm. Participants were asked to choose the least polar version. They were also asked questions about whether the three versions were about the same topic, and whether three versions were semantically similar. The participants were also asked to rate the depolarization success of the semi-automated and fully automated versions. All questions were measured on a 7-point scale.

Measurements

Demographic measures included gender, racial category, years of education, and age. Political attitude was measured by a 7-point scale (1-Extremely liberal; 4-Moderate; 7-Extremely conservative) as:

- “Generally speaking, where would you place yourself on the following scale?”

Topics	Mean (SD) of the Extent of Polarity		
	Original	Semi-auto	Fully-auto
Social Issues	4.77 (1.48)	4.22 (1.54)	4.03 (1.69)
Immigration	4.63 (2.04)	4.63 (1.78)	4.4 (1.72)
Miscellaneous	5.5 (1.65)	4.9 (1.92)	5.22 (1.73)
Healthcare	3.95 (1.69)	3.67 (1.42)	3.4 (1.66)
Climate	4.07 (1.48)	3.63 (1.44)	3.55 (1.55)
Trade	4.87 (1.86)	4.53 (1.67)	4.4 (1.59)
Election	4.28 (1.93)	3.98 (1.89)	4.13 (1.75)
Foreign Policy	4.8 (1.42)	4.48 (1.56)	4.25 (1.76)
Education	4.38 (1.93)	4.47 (1.91)	4.43 (1.8)
Criminal Justice	3.63 (1.77)	3.37 (1.58)	3.55 (1.52)
Vaccination	4.13 (1.73)	4.05 (1.66)	3.88 (1.59)

Table 4: Mean of the extent of polarity for each topic in three modes (original, semi-automatic and fully-automatic). We also mark the SD (standard deviation) for each mean.

Party affiliation was measured by 7-point scale (1-Strong Democratic; 4-Independent; 7-Strong Republican) as:

- “Generally speaking, do you usually think of yourself as a Republican, a Democrat, or an Independent?”

After reading each paragraph without knowing the version type (i.e., whether it was an original article or a depolarized article), participants were asked about the bias, ideology, and the readability of the paragraph. Bias was measured by as (1- Extremely unbiased to 7- Extremely biased):

- “How biased do you think this story is?”

Ideology was measured as (1-Extremely liberal; 4-Moderate; 7-Extremely conservative):

- “What ideology do you think this story is?”

The readability measure included five items adapted from a previous study by Haim et al. (Haim and Graefe 2017), namely, well-written, concise, comprehensive, coherent, and clear. The reliability of the five items of readability was $Cronbach's \alpha = 0.94$, $M = 19.51$, $SD = 8.47$, which was very high. Thus, we can average five metrics into one as a general reliability index.

In the post-exposure questionnaire, three questions were used to test the content preserving performance:

- “Do you agree that the three paragraphs above talk about the same topic?”
- “Do you agree that the three paragraphs above hold the same political views?”
- “Do you agree that the three paragraphs above are semantically similar?”

Results

Depolarization. As mentioned earlier, the extent of polarity was measured on a 7-point scale, with higher values corresponding to greater polarity. The average polarity of the

Topics	Original vs. Semi			Original vs. Fully		
	<i>t</i>	<i>df</i>	<i>p-Value</i>	<i>t</i>	<i>df</i>	<i>p-Value</i>
Social Issues	2.88	59	0.005 **	3.69	59	0.00 ***
Immigration	0.00	59	1.00	1.37	59	0.18
Miscellaneous	3.49	59	0.00 ***	1.29	59	0.20
Healthcare	1.25	59	0.22	2.45	59	0.017 *
Climate	2.39	59	0.02 **	2.47	59	0.016 *
Trade	2.88	59	0.005 **	3.5	59	0.00 ***
Election	1.8	59	0.08	0.76	59	0.45
Foreign Policy	2.45	59	0.017 *	2.74	59	0.008 **
Education	-0.4	59	0.69	-0.23	59	0.82
Criminal Justice	1.49	59	0.14	0.34	59	0.74
Vaccination	0.52	59	0.61	1.2	59	0.23

Table 5: Two paired samples t-tests on the polarity of paired paragraphs from 11 topics. The purpose is to determine whether there is statistical difference in polarity between the original text and the output of the depolarized-text generation algorithm in both the semi-automatic and fully-automatic modes. (* corresponds to $p < 0.05$, ** to $p < 0.01$ and *** to $p < 0.001$.)

articles (as specified by the participants) for each version of the stories is shown in Table 4. Furthermore, we used paired sample t-tests to examine whether there were significant differences in perceived polarity between the original and depolarized stories. These results are shown in Table 5. The results show that for certain topics, there were statistically significant reductions in participants’ perceived polarity when comparing original articles to their depolarized versions (for both modes). (The magnitude of the reduction can be seen in Table 4 and the statistical significance in Table 5. In general, there were no notable differences between the semi-automatic and fully-automatic versions, which means that further improvements to the framework need to be focused on the attribute-aware word embeddings (which are used by both the semi-automatic and fully-automatic versions) our depolarization algorithm could achieve comparable and competitive results as human-aided method.

Readability (Pairwise Comparison)				
Pair	Mean (SD)	<i>t</i>	<i>df</i>	<i>p-Value</i>
Original vs. Semi	4 (1.67)	0.891	50	0.377
	3.93 (1.69)			
Original vs. Fully	4 (1.67)	-0.558	50	0.579
	4.07 (1.71)			
Semi vs. Fully	3.93 (1.69)	-1.094	50	0.279
	4.07 (1.71)			

Table 6: T-test on the readability of original text, the generation from semi-automatic mode and from fully-automatic mode. We do not observe

Readability. Our results show that our depolarization algorithms, TADA, did not weaken readability. Paired sample t-tests showed that no significant difference existed among readability of the three versions (Table 6). The semi-automatic and fully-automatic edited versions were as readable as the original text.

Content Preserving. Three questions were used to test the content preserving performance, as listed in the “Measurements” section. Recall that these questions were asked post-exposure where participants were shown the three versions of the texts at the same time and asked three questions. The first question was intended as a soft content preserving check, since it is possible that the topic of a text remains the same while the political position is completely flipped. The second questions is a harder check as it is asking about the political views contained in the text; it is again easy to imagine a scenario where the political views are preserved while the meaning is changed. Finally, the third question gets at the preservation of the meaning (semantics) conveyed in the text. Results showed that on average 81.2% of the participants agreed that the three versions of the texts were talking about the same topic, 75.2% of the participants agreed that the three versions of the texts had the same political view, and 79% of the participants agreed that the three versions of the texts had the same meaning. We can see that across all three measures, over three quarters of the participants felt that the content of the texts had been preserved in the depolarization process.

Conclusions

In this work, we proposed a framework for depolarizing news articles which can be used either fully-automatically or semi-automatically. At the core of our framework is a novel attribute-aware word embedding model which is trained on two attributes (ideology and topic), using data from Media Cloud. The model can be used to detect polar content in articles. We also proposed a probabilistic algorithm, called TADA, that can depolarize text using our attribute-aware word embedding model. We evaluated our framework and its components empirically, qualitatively, and through human evaluations, showing that our framework is capable of depolarizing articles for certain topics while preserving the meaning and without damaging readability.

While our best-tested model performs well, there is still much room for improvement. For example, the polarity detection of our model is not truly contextualized, we need to further consider the influence of context on the polarity score. Additionally, since as part of our attribute-aware model we create separate word embeddings for each attribute, our model cannot be easily scaled to handle a large number of attributes. The limitation of memory could be a potential bottleneck for the extension. Future work could involve employing contextual language models in our framework, more memory efficient methods for creating multi-attribute-aware embeddings and exploring continuous attributes. Moreover, the general field that this paper falls under suffers from a lack of established evaluation benchmarks and metrics. More work needs to be done to address these shortcomings and unify the field.

References

- Abadi, M. 2017. Democrats and republicans speak different languages and it helps explain why we're so divided. *Business Insider*.
- Alzantot, M.; Sharma, Y.; Elgohary, A.; Ho, B.-J.; Srivastava, M.; and Chang, K.-W. 2018. Generating natural language adversarial examples. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2890–2896.
- Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2016. Enriching word vectors with subword information. *arXiv preprint arXiv:1607.04606*.
- Bowman, S. R.; Vilnis, L.; Vinyals, O.; Dai, A. M.; Jozefowicz, R.; and Bengio, S. 2015. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*.
- Buhrmester, M.; Kwang, T.; and Gosling, S. D. 2011. Amazon's mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on psychological science* 6(1):3–5.
- Chen, X.; Shi, Z.; Qiu, X.; and Huang, X. 2017. Adversarial multi-criteria learning for Chinese word segmentation. *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)* 1:1193–1203.
- Chen, L.; Dai, S.; Tao, C.; Shen, D.; Gan, Z.; Zhang, H.; Zhang, Y.; and Carin, L. 2018a. Adversarial text generation via feature-mover's distance. *Advances in Neural Information Processing Systems 2018-Decem(NeurIPS)*:4666–4677.
- Chen, W.-F.; Wachsmuth, H.; Al Khatib, K.; and Stein, B. 2018b. Learning to flip the bias of news headlines. In *Proceedings of the 11th International Conference on Natural Language Generation*, 79–88.
- Dai, N.; Liang, J.; Qiu, X.; and Huang, X. 2019a. Style transformer: Unpaired text style transfer without disentangled latent representation. *arXiv preprint arXiv:1905.05621*.
- Dai, N.; Liang, J.; Qiu, X.; and Huang, X. 2019b. Style transformer: Unpaired text style transfer without disentangled latent representation. *arXiv preprint arXiv:1905.05621*.
- Dardis, F. E.; Baumgartner, F. R.; Boydston, A. E.; De Boef, S.; and Shen, F. 2008. Media framing of capital punishment and its impact on individuals' cognitive responses. *Mass Communication & Society* 11(2):115–140.
- Gentzkow, M., and Shapiro, J. M. 2010. What drives media slant? evidence from us daily newspapers. *Econometrica* 78(1):35–71.
- Gillani, N., and Levy, R. 2019. Simple dynamic word embeddings for mapping perceptions in the public sphere. *arXiv.org*.
- Glaeser, E. L., and Goldin, C. 2007. *Corruption and Reform: Lessons from America's Economic History*. University of Chicago Press.
- Gong, H.; Bhat, S.; Wu, L.; Xiong, J.; and Hwu, W.-m. 2019. Reinforcement learning based text style transfer without parallel training corpus. *arXiv preprint arXiv:1903.10671*.
- Groseclose, T., and Milyo, J. 2005. A measure of media bias. *The Quarterly Journal of Economics* 120(4):1191–1237.
- Guu, K.; Hashimoto, T. B.; Oren, Y.; and Liang, P. 2018. Generating sentences by editing prototypes. *Transactions of the Association for Computational Linguistics* 6:437–450.
- Haim, M., and Graefe, A. 2017. Automated news: Better than expected? *Digital journalism* 5(8):1044–1059.
- Huang, S.; Wu, Y.; Wei, F.; and Luan, Z. 2019. Dictionary-Guided Editing Networks for Paraphrase Generation. *Proceedings of the AAAI Conference on Artificial Intelligence* 33:6546–6553.
- Huillery, E. 2009. History matters: The long-term impact of colonial public investments in french west africa. *American economic journal: applied economics* 1(2):176–215.
- Iyyer, M.; Enns, P.; Boyd-Graber, J.; and Resnik, P. 2014. Political ideology detection using recursive neural networks. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1113–1122.
- John, V.; Mou, L.; Bahuleyan, H.; and Vechtomova, O. 2018. Disentangled representation learning for non-parallel text style transfer. *arXiv preprint arXiv:1808.04339*.
- Kahneman, D., and Tversky, A. 2013. Choices, values, and frames. In *Handbook of the fundamentals of financial decision making: Part I*. World Scientific. 269–278.
- Lample, G.; Subramanian, S.; Smith, E. M.; Denoyer, L.; Ranzato, M. A.; and Boureau, Y.-L. 2019. Multiple-Attribute Text Rewriting. *International Conference on Learning Representations* 1–20.
- Lan, Z.; Chen, M.; Goodman, S.; Gimpel, K.; Sharma, P.; and Soricut, R. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Lemmon, T. G. 2013. Washington vs. d.c.: How the shutdown pits the government against its city. *The Atlantic*.
- Li, J.; Jia, R.; He, H.; and Liang, P. 2018. Delete, retrieve, generate: A simple approach to sentiment and style transfer. *arXiv preprint arXiv:1804.06437*.
- Lin, Y.-R.; Bagrow, J.; and Lazer, D. 2011. More voices than ever? quantifying media bias in networks. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 5.
- Liu, R.; Jia, C.; Wei, J.; Xu, G.; Wang, L.; and Vosoughi, S. 2021. Mitigating political bias in language models through reinforced calibration. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Liu, R.; Jia, C.; and Vosoughi, S. 2021. A transformer-based framework for neutralizing and reversing the political polarity of news articles. *Proceedings of the ACM on Human-Computer Interaction*.
- Logeswaran, L.; Lee, H.; and Bengio, S. 2018. Content preserving text generation with attribute controls. *Advances in Neural Information Processing Systems 2018-Decem(Nips)*:5103–5113.

- Luntz, F. 2007. *Words that work: It's not what you say, it's what people hear*. Hachette UK.
- Marangos, J. 2008. The evolution of the anti-washington consensus debate: From 'post-washington consensus' to 'after the washington consensus'. *Competition & Change* 12(3):227–244.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119.
- Miller, Z. 2011. Revealed: How to talk like a republican. *Business Insider*.
- Mueller, J.; Gifford, D.; and Jaakkola, T. 2017. Sequence to better sequence: Continuous revision of combinatorial structures. *34th International Conference on Machine Learning, ICML 2017* 5(1):3900–3916.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, 311–318. Association for Computational Linguistics.
- Prabhumoye, S.; Tsvetkov, Y.; Salakhutdinov, R.; and Black, A. W. 2018. Style transfer through back-translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 866–876.
- Ren, S.; Deng, Y.; He, K.; and Che, W. 2019. Generating natural language adversarial examples through probability weighted word saliency. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 1085–1097.
- Scammell, M., and Semetko, H. 2018. *The media, journalism and democracy*. Routledge.
- Schnabel, T.; Labutov, I.; Mimno, D.; and Joachims, T. 2015. Evaluation methods for unsupervised word embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 298–307.
- Shen, T.; Lei, T.; Barzilay, R.; and Jaakkola, T. 2017. Style transfer from non-parallel text by cross-alignment. In *Advances in neural information processing systems*, 6830–6841.
- Shoemark, P.; Sur, D.; Shrimpton, L.; Murray, I.; and Goldwater, S. 2017. Aye or naw, whit dae ye hink? scottish independence and linguistic identity on social media. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, 1239–1248.
- Singh, A., and Palod, R. 2018. Sentiment transfer using seq2seq adversarial autoencoders. *arXiv preprint arXiv:1804.04003*.
- Subramanian, S.; Lample, G.; Smith, E. M.; Denoyer, L.; Ranzato, M.; and Boureau, Y.-L. 2018. Multiple-attribute text style transfer. *arXiv preprint arXiv:1811.00552*.
- Tian, Y.; Hu, Z.; and Yu, Z. 2018. Structured content preservation for unsupervised text style transfer. *arXiv preprint arXiv:1810.06526*.
- Wu, C.; Ren, X.; Luo, F.; and Sun, X. 2019. A hierarchical reinforced sequence operation method for unsupervised text style transfer. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4873–4883.
- Xu, J.; Sun, X.; Zeng, Q.; Ren, X.; Zhang, X.; Wang, H.; and Li, W. 2018. Unpaired sentiment-to-sentiment translation: A cycled reinforcement learning approach. *arXiv preprint arXiv:1805.05181*.
- Yang, Y., and Eisenstein, J. 2017. Overcoming language variation in sentiment analysis with social attention. *Transactions of the Association for Computational Linguistics* 5:295–307.
- Yano, T.; Resnik, P.; and Smith, N. A. 2010. Shedding (a thousand points of) light on biased language. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, 152–158. Association for Computational Linguistics.
- Youngmann, B. 2019. A simple approach to remove bias from text. *Technical Report*.
- Zhao, J.; Kim, Y.; Zhang, K.; Rush, A. M.; and LeCun, Y. 2018. Adversarially regularized autoencoders. *35th International Conference on Machine Learning, ICML 2018* 13:9405–9420.
- Zhou, C., and Neubig, G. 2017. Multi-space variational encoder-decoders for semi-supervised labeled sequence transduction. In *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, volume 1, 310–320.