

# Encoding Social Information with Graph Convolutional Networks for Political Perspective Detection in News Media

**Chang Li**  
Purdue University  
li1873@purdue.edu

**Dan Goldwasser**  
Purdue University  
dgoldwas@purdue.edu

## Abstract

Identifying the political perspective shaping the way news events are discussed in the media is an important and challenging task. In this paper, we highlight the importance of contextualizing social information, capturing how this information is disseminated in social networks. We use Graph Convolutional Network, a recently proposed neural architecture for representing relational information, to capture this information and show that even with little social information classification can be significantly improved.

## 1 Introduction

Over the last decade we witness a dramatic change in the way information is generated and disseminated. Instead of a few dedicated sources that employ reporters and fact checkers to ensure the validity of the information they provide, social platforms now provide the means for any user to distribute their content, resulting in a sharp increase in the number of information outlets and articles covering news events. As a direct result of this process, the information provided is often shaped by their underlying perspectives, interests and ideologies. For example, consider the following two snippets discussing the comments made by a Democratic Senator regarding the recent U.S. government shutdown.

**thehill.com** (*Center*)

Sen. Mark Warner (D-Va.) on Sunday blasted President Trump for his “inept negotiation” to bring an end to the ongoing partial government shutdown. Warner, the ranking member of the Senate Intelligence Committee, lamented the effect the shutdown has had on hundreds of thousands of federal workers who have been furloughed or forced to work without pay.

**infowars.com** (*Right*)

Senator Mark Warner (D-Va.) is being called out on social media for his statement on the partial government shutdown. Warner blamed the “suffering” of federal workers and contractors on President Trump in a Sunday tweet framing Trump as an “inept negotiator”. Twitter users pointed out that Democrats are attending a Puerto Rican retreat with over 100 lobbyists and corporate executives.

Despite the fact that both articles discuss the same event, they take very different perspectives. The first reporting directly about the comments made, while the second one focuses on negative reactions to these comments. Identifying the perspective difference and making it explicit can help strengthen trust in the newly-formed information landscape and ensure that all perspectives are represented. It can also help lay the foundation for the automatic detection of false content and rumors and help identify information echo-chambers in which only a single perspective is highlighted.

Traditionally, identifying the author’s perspective is studied as a text-categorization problem (Greene and Resnik, 2009; Klebanov et al., 2010; Recasens et al., 2013a; Iyyer et al., 2014a; Johnson and Goldwasser, 2016), focusing on linguistic indicators of bias or issue-framing phrases indicating their authors’ bias. These indicators can effectively capture bias in ideologically-charged texts, such as policy documents or political debates, which do not try to hide their political leaning and use a topic-focused vocabulary. Identifying the authors’ bias in news narratives can be more challenging. News articles, by their nature, cover a very large number of topics resulting in a diverse and dynamic vocabulary that is continuously updated as new events unfold. Furthermore, unlike purely political texts, news narratives attempt to maintain credibility and seem impartial. As a result, bias is introduced in subtle ways, usu-

ally by emphasizing different aspects of the story.

Our main insight in this paper is that the social context through which the information is propagated can be leveraged to alleviate the problem, by providing both a better representation for it, and when direct supervision is not available, a distant-supervision source based on information about users who endorse the textual content and spread it. Several recent works dealing with information dissemination analysis on social networks, focused on analyzing the interactions between news sources and users in social networks (Volkova et al., 2017; Glenski et al., 2018; Ribeiro et al., 2018). However, given the dynamic, and often adversarial setting of this domain, the true source of the news article might be hidden, unknown or masked by taking a different identity. Instead of analyzing the documents’ sources, our focus is to use social information, capturing how information is shared in the network, to help guide the text representation and provide additional support when making decisions over textual content.

We construct a *socially-infused textual representation*, by embedding in a single space the news articles and the social circles in which these articles are shared so that the political biases associated with them can be predicted. Figure 1 describes these settings. The graph connects article nodes via activity-links to users nodes (*share*), and these users in turn are connected via social-links (*follow*) to politically affiliated users (e.g., the Republican or Democratic parties twitter accounts). We define an embedding objective capturing this information, by aligning the documents representation, based on content, with the representation of users who share these documents, based on their social relations. We use a recently proposed graph embedding framework, Graph Convolutional Networks (GCN) (Kipf and Welling, 2016, 2017) to capture these relationships. GCNs are neural nets operating on graphs, and similar to LSTMs on sequences, they create node embeddings based on the graph neighborhood of a given node. In the context of our problem, the embedding of a document takes into account the textual content, but also the social context of users who share it, and their relationships with other users with known political affiliations. We compare this powerful approach with traditional graph embedding methods that capture local relationships between nodes.

Given the difficulty of providing direct super-

vision in this highly dynamic domain, we study this problem both when direct supervision over the documents is available, and when using *distant-supervision*, in which the document level classification depends on propagating political tendencies through social network, which is often incomplete and provides conflicting information.

To study these settings we focus on U.S. news coverage. Our corpus consists of over 10k articles, covering more than 2k different news events, about 94 different topics, taking place over a period of 8 years. We remove any information about the source of the article (both meta-data and in the text) and rely only on the text and the reactions to it on social media. To capture this information, we collected a set of 1.6k users who share the news articles on Twitter and a handful of politically-affiliated users followed by the sharing users, which provide the distant supervision. We cast the problem as a 3-class prediction problem, capturing left-leaning bias, right-leaning bias or no bias (center).

Our experimental results demonstrate the strength of our approach. We compare direct text classification or node classification methods to our embedding-based approach in both the fully supervised and distant supervised settings, showing the importance of socially infused representations.




Figure 1: Information Flow Graph

## 2 Related Work

The problem of perspective identification is typically studied as a supervised learning task (Lin et al., 2006; Greene and Resnik, 2009), in which

a classifier is trained to differentiate between two specific perspectives. For example, the *bitterlemons* dataset consisting of 594 documents describing the Israeli and Palestinian perspectives. More recently, in SemEval-2019, a hyperpartisan news article detection task was suggested<sup>1</sup>. The current reported results on their dataset are comparable to ours, when using text information alone, demonstrating that it is indeed a challenging task. Other works use linguistic indicators of bias and expressions of implicit sentiment (Greene and Resnik, 2009; Recasens et al., 2013a; Choi and Wiebe, 2014; Elfardy et al., 2015). In recent years several works looked at indications of framing bias in news articles (Baumer et al., 2015; Budak et al., 2016; Card et al., 2016; Field et al., 2018; Morstatter et al., 2018). We build on these work to help shape our text representation approach.

Recent works looked at false content identification (Volkova et al., 2017; Patwari et al., 2017), including a recent challenge<sup>2</sup> identifying the relationship between an article’s title and its body. Unlike these, we do not assume the content is false, instead we ask if it reflects a different perspective.

Using social information when learning text representations was studied in the context of graph embedding (Pan et al., 2016), extending traditional approaches that rely on graph relations alone (Perozzi et al., 2014; Tang et al., 2015a; Grover and Leskovec, 2016) and information extraction and sentiment tasks (Yang et al., 2016a; West et al., 2014). In this work we focus on GCNs (Kipf and Welling, 2017; Schlichtkrull et al., 2018), a recent framework for representing relational data, that adapts the idea of convolutional networks to graphs. Distant supervision for NLP tasks typically relies on using knowledge-bases (Mintz et al., 2009), unlike our setting that uses social information. Using user activity and *known* user biases was explored in (Zhou et al., 2011), our settings are far more challenging as we do not have access to this information.

### 3 Dataset Description

We collected 10,385 news articles from two news aggregation websites<sup>3</sup> on 2,020 different events discussing 94 event types, such as elections, terrorism, etc. The websites provide news cover-

Articles	10,385	Twitter Users	1,604
-Left	3,931	Pol. Users	135
-Right	2,290	Left Pol. Users	49
-Center	4,164	Right Pol. Users	51
Sources	86	Center Pol. Users	35
Types	94	Avg # shared per Article	23.29
Events	2,020	Avg # pol. users followed	20.36

Table 1: Dataset Statistics

age from multiple perspectives, indicating the bias of each article using crowdsourced and editorial reviewed approaches<sup>4</sup>. We preprocessed all the documents to remove any information about the source of the article.

We collected social information consisting of Twitter users who share links to the collected articles. We focused on Twitter users who follow political users and share news articles frequently (100 articles minimum). We found 1,604 such Twitter users. The list of political users was created by collecting information about active politically affiliated users. It consists of 135 Twitter users who are mainly politicians, political journalists and political organizations. The set of political users and Twitter users are disjoint. The summary of the dataset is shown in Table 1.

**Data Folds** We created several data splits to evaluate our model in the supervised settings, based on three criteria: *randomly* separated, *event* separated and *time* separated splits. In the event-separated case, we divide the news articles such that all articles covering the same news event will appear in a single fold. For the time-separated case, we sort the publication dates (from oldest to latest) and divide them in three folds. Each time one fold is used as training data (33%) and other two combined as test data (66%). We use the same folds throughout the experiment of supervised classification for evaluation purpose.

#### Constructing the Social Information Graph

We representing the relevant relationships as an information graph, similar to the one depicted in Figure 1. The social information graph  $G = \{V, E\}$ , consisting of several different types of vertices and edges, is defined as follows:

- Let  $P \subset V$  denote the set of the *political users*. These are Twitter users with a clear, self-reported, political bias. They may be the accounts of politicians (e.g., Sarah Palin, Nancy

<sup>1</sup><https://webis.de/events/semEval-19/>

<sup>2</sup><http://www.fakenewschallenge.org>

<sup>3</sup>[Memeorandum.com](http://Memeorandum.com) and [Allsides.com](http://Allsides.com)

<sup>4</sup><https://www.allsides.com/media-bias/media-bias-rating-methods>

Pelosi), political writers in leading newspapers (e.g., Anderson Cooper) or political organizations (e.g., GOP, House Democrats). Note that even political users that share a general political ideology can differ significantly in the type of issues and agenda they would pursue, which would be reflected in their followers.

- Let  $U \subset V$  denote the set of *Twitter users* that actively spread content by sharing news articles. The political bias of these users is not directly known, only indicated indirectly through the *political users* they follow on Twitter.
- Let  $A \subset V$  denote the set of news articles shared by the *Twitter users* ( $U$ ).

The graph vertices are connected via a set of edges described hierarchically, as follows:

- $E_{UP} \subset E$ : All the Twitter users are connected to the political users whom they follow. Note that a Twitter user may be connected to many different political users.
- $E_{AU} \subset E$ : All the articles are connected to the Twitter users who share them. Note that an article may be shared by many different Twitter users.

## 4 Text and Graph Model

Our goal is to classify news articles into 3-classes corresponding to their bias. Since we have both the textual and social information for the news articles. We can obtain representations of them using either text or graph models. In this section, we briefly go through the text representation methods, and then move to describe the graph based models we considered in this paper.

### 4.1 Text Representations and Linguistic Bias Indicators

To predict the bias of the news articles, we can consider it as a document classification task. We use the textual content of a news article to generate feature representation. Deciding on the appropriate representation for this content is one of the key design choices. Previous works either use traditional, manually engineered representations for capturing bias (Recasens et al., 2013b) or use latent representations learned using deep learning methods (Iyyer et al., 2014b). We experimented with several different choices of the two alternatives, and compared them by training a classifier

for bias prediction over the document directly. The results of this experiments are summarized in Table 2. Due to space constraints, we provide a brief overview of these alternatives, and point to the full description in the relevant papers.

**Linear BoW** Unigram features were used. The articles consist of 77,772 unique tokens. We used TFIDF vectors as unigram features obtained by using scikit-learn (Pedregosa et al., 2011).

**Bias Features** These features are content based features drawn from a wide range of literature on political bias, persuasion, and misinformation, capturing structure, sentiment, topic, complexity, bias and morality in the text. We used the resources in (Horne et al., 2018b) to generate 141 features based on the news article text. It is shown in (Horne et al., 2018a) that these features can be used in the binary hyper-partisan task effectively.

**Averaged Word Embedding (WE)** An averaged vector of all the documents words using the pre-trained GloVe word embeddings (Pennington et al., 2014) were used to represent the entire article. This is the basic way to use the pre-trained word embeddings.

**Skip-Thought Embedding** Unlike the Averaged word vector that does not capture context, we also used a sentence level encoder, Skip-Thought (Kiros et al., 2015), to generate text representations. We regard each document as a long sentence and map it directly to a 4800-dimension vector.

### Hierarchical LSTM over tokens and sentences

We used a simplified version of the Hierarchical LSTM model (Yang et al., 2016b). In this case documents are first tokenized into sentences, then each sentence was tokenized into words. We used a word-level LSTM to construct a vector representation for each sentence, by taking the average of all the hidden states. Then, we ran another single layer unidirectional LSTM over the sentence representations to get the document representation by taking average of all the hidden states.

## 4.2 Graph-Based Representations

In addition to the textual information, the news articles are also part of the information network defined in Section 3. It is clear that news articles shared by the same Twitter users are likely to have the same bias, and users who share a lot of



news in common are close in their political preference. It is similar for the relation between sharing users and political users. Therefore bias can be predicted given the social information of a news article as well. We design our embedding function to map all graph nodes into a low dimensional vector space, such that the graph relationships are preserved in the embedding space. In the shared embedding space, nodes that are connected (or close) in the graph should have higher similarity scores between their vector representations.

#### 4.2.1 Directly Observed Relationships in Graph (DOR)

Our first embedding approach aims to preserve the local pairwise proximity between two vertices directly. This is similar to first-order graph embedding methods (Tang et al., 2015b). There are two different relations observed in the graph: Twitter user to political user (follow) and news article to Twitter user (share). We construct our embedding over multiple views of the data, each view  $V$  corresponds to a specific type of graph relation. We can then define an loss function  $L_V$  for each view  $V$  as follows:

- Twitter User to Political User (UP): This objective maximizes the similarity of a Twitter user,  $u$  and all the political users in the set  $P_u \subset P$ , where  $P_u$  is the set of political users that  $u$  follows.

$$L_{UP} = - \sum_{u \in U} \sum_{p \in P_u} \log P(p|u) \quad (1)$$

- News Article to Twitter User (AU): This objective maximizes the similarity of a news articles,  $a$  and all the Twitter users in the set  $U_a \subset U$ , where  $U_a$  is the set of Twitter users who shared news article  $a$  on Twitter.

$$L_{AU} = - \sum_{a \in A} \sum_{u \in U_a} \log P(u|a) \quad (2)$$

All the conditional probabilities can be computed using a softmax function. Taking  $P(p|u)$  as an example:

$$P(p|u) = \frac{\exp(e_u^T e_p)}{\sum_{q \in P} \exp(e_u^T e_q)} \quad (3)$$

Computing Eq. 1 and Eq. 2 can be expensive due to the size of the network. To address this

problem, we refer to the popular negative sampling approach (Mikolov et al., 2013), which reduce the time complexity to be proportional to the number of positive example pairs (i.e. number of edges in our case).

The loss defined for the two views are summed with the classification loss defined in Eq. 9 as the final loss function to be optimized in DOR embedding model.

$$L_{DOR} = L_{clf} + L_{UP} + L_{AU} \quad (4)$$

#### 4.2.2 Graph Convolutional Networks (GCN)

Graph Convolutional Networks is an efficient variant of convolutional neural networks which operate directly on graphs. It can be regarded as special cases of a simple differentiable message-passing framework (Gilmer et al., 2017):

$$h_i^{(l+1)} = \sigma \left( \sum_{j \in N(i)} M^{(l)}(h_i^{(l)}, h_j^{(l)}) \right) \quad (5)$$

where  $h_i^{(l)} \in \mathbb{R}^{d^{(l)}}$  is the hidden state of node  $v_i$  in the  $l$ -th layer of the neural network, with  $d^{(l)}$  as the dimensionality of representation at layer  $l$ .  $N(i)$  is the set of direct neighbors of node  $v_i$  (usually also include itself). Incoming messages from the local neighborhood are aggregated together and passed through the activation function  $\sigma(\cdot)$ , such as  $\tanh(\cdot)$ .  $M^{(l)}$  is typically chosen to be a (layer-specific) neural network function. Kipf and Welling (2017) used a simple linear transformation  $M^{(l)}(h_i^t, h_j^t) = W^{(l)} h_j$  where  $W^{(l)}$  is a layer-specific weight matrix.

This linear transformation has been shown to propagate information effectively on graphs. It leads to significant improvements in node classification (Kipf and Welling, 2017), link prediction (Kipf and Welling, 2016), and graph classification (Duvenaud et al., 2015).

One GCN layer can be expressed as follows:

$$H^{(l+1)} = \sigma(\hat{A}H^{(l)}W^{(l)}) \quad (6)$$

where  $\hat{A}$  is the normalized adjacency matrix, and  $W^{(l)}$  is the layer-specific trainable weight matrix.  $H^{(l)} \in \mathbb{R}^{N \times D^{(l)}}$  is the matrix of hidden states in the  $l$ -th layer.  $H^{(0)} = X$  is the input vectors. It can either be one-hot representations of nodes or features of the nodes if available.  $\sigma(\cdot)$  is the activation function.

Multiple GCN layers can be stacked in order to capture high-order relations in the graph. We consider a two-layer GCN in this paper for semi-supervised node classification. Our forward model takes the form:

$$V = \text{softmax}\left(\hat{A} \tanh\left(\hat{A}XW^{(0)}\right)W^{(1)}\right) \quad (7)$$

where  $X$  is the input matrix with one-hot representations and  $V$  is the representation matrix for all nodes in the graph.

Figure 2 shows an example of how our GCN model aggregate information from a node’s local neighborhood. The orange document is the node of interest. Blue edges link to first order neighbors and green edges link to second order neighbors.




Figure 2: Example of Unfolding of GCN Computational Graph

### 4.3 Document Classification

The representation  $v$  of an news article (obtained with text models or graph models) capture the high level information of the document. It can be used as features for predicting bias label with a feed-forward network.

$$p = \text{softmax}(W_c v + b_c) \quad (8)$$

We use the negative log likelihood of the correct labels as classification training loss:

$$L_{clf} = - \sum_a \log p_{a,j} \quad (9)$$

where  $j$  is the bias label of news article  $a$ .

## 5 Joint Model

Given that we have the two parts of information available for the news articles, namely the textual

and social one, it is natural to make the prediction combining both of them. We propose to align the representations of the same document from graph and text models in a joint training fashion. The objective function for the alignment is:

$$L_{align} = - \sum_{a \in A} \log P(a^g | a^t) \quad (10)$$

where  $a^t$  is the embedding for document  $a$  based on its content, and  $a^g$  is the embedding for document  $a$  based on graph structures.  $P(a^g | a^t)$  is defined the same way as in Eq. 3. Negative sampling is again utilized to reduce time complexity. By connecting the text and graph embedding of the same news articles, it allows bias signal to flow between these two sides. Therefore text model may learn from the social signal and graph model may use textual content to adjust the embeddings as well. We describe the loss function for the joint model in full supervision and distant supervision cases.

**Full Supervision** In the full supervision case, the loss consists of three parts, namely classification loss of text model ( $L_{clf}^T$ ), classification loss of graph model ( $L_{clf}^G$ ), and the loss to align embeddings from text and graph models ( $L_{align}$ ).

$$L_{joint} = \alpha L_{clf}^T + \beta L_{clf}^G + \gamma L_{align} \quad (11)$$

Here  $\alpha$ ,  $\beta$  and  $\gamma$  are hyper-parameters to adjust the contribution of the three parts. We set all of them to default value 1 in experiments in this paper.

**Distant Supervision** Unlike the full supervision case where we have training labels for documents, we only have access to labels of political users. However, we can still make use of the joint model by regarding the labels predicted by the graph model as supervision on news articles. The assumption is the joint training can still benefit from the noisy training labels on documents.

**Inference** Now that we have multiple representations for the news articles, we define a simplified inference process, which also takes into account the users sharing the document. At test time, we can make prediction of a news article with the embeddings from text model (Text), the embeddings from graph model (Graph), and the embeddings of sharing users who shared this article (User). The last method (User) works by averaging bias prediction scores  $s_u^b$  for all Twitter users shared arti-

cle  $a$ . Bias prediction scores is computed in Eq. 8 before the softmax( $\cdot$ ) applied.

$$\arg \max_b \frac{\sum_{u \in U_a} s_u^b}{|U_a|} \quad (12)$$

Finally a combination of two or three of the scores listed above can be added to make the decision.

## 6 Experiments

We designed our experiments to evaluate the contribution of social information in both the fully supervised setting, and when only distant supervision is available through social graph. We begin by evaluating several text classification models that help contextualize the social information. Finally, we evaluate our model’s ability to make predictions when very little social information is available for test news articles.

### 6.1 Implementation Details

We used the spaCy toolkit for preprocessing the documents. All models are implemented with PyTorch (Paszke et al., 2017)<sup>5</sup>. Hyperbolic tangent (tanh) is used as non-linear activation function. We use feed-forward neural network with one hidden layer for the bias prediction task given textual or social representation. The sizes of LSTM hidden states for both word level and sentence level are 64. The sizes of hidden states for both GCN layers are 16. For the training of the neural network, we used Adam optimizer (Kingma and Ba, 2014) to update parameters. We use 5% of training data as validation set. We run the training for 200 epochs (50 epochs for models using HLSTM), and select the best model based on performance on validation set. Other parameters in our model includes negative sample size  $k=5$ , mini-batch size  $b=30$  (mini-batch update only used for models with HLSTM).

### 6.2 Experimental Results

**Text Classification Results** The result of supervised text classification is summarized in Table 2. We report the accuracy of bias prediction. It is clearly shown that HLSTM outperforms the other methods in supervised text classification setting. Also, adding the hand engineered bias features with HLSTM representation does not help to improve performance.

Model	Split	Text
Majority	Rand	40.10
	Event	40.10
	Time	40.50
Linear BoW	Rand	58.47
	Event	59.88
	Time	55.41
Bias Feat.	Rand	54.06
	Event	53.51
	Time	52.96
Avg WE	Rand	59.37
	Event	59.37
	Time	53.46
SkipThought	Rand	68.67
	Event	66.35
	Time	60.89
HLSTM	Rand	<b>74.59</b>
	Event	<b>73.55</b>
	Time	<b>66.98</b>
HLSTM + Bias Feat.	Rand	69.32
	Event	69.87
	Time	66.79

Table 2: Supervised Classification Using Textual Features

**Network Classification Results** We show the results predicting bias using graph information alone, without text, in Table 3. The GCN model outperforms DOR significantly in each of the four settings. Similar with the results of text classification, performance on random and event splits are comparable. However, there is a sharp drop in performance for time split. This can be explained by the fact that temporally separated news event will discuss different entities and world events and as a result will have very different word distributions. Event-separated splits are less susceptible to this problem, as similar figures and topics are likely to be discussed in different events.

Model	Split	Graph	User	G+U
DOR	Rand	74.74	72.02	74.57
	Event	74.87	72.74	75.18
	Time	65.65	65.07	65.36
	Dist	56.45	56.95	56.54
GCN	Rand	88.65	78.83	<b>88.89</b>
	Event	<b>88.78</b>	76.11	88.70
	Time	81.14	71.31	<b>82.00</b>
	Dist	63.72	40.08	<b>67.03</b>

Table 3: Classification Results Using Social Relations in Full Supervised and Distant Supervised Setting

**Joint Model Results** Table 4 shows the results of our joint model. By aligning the text and graph embeddings in the joint training, both of them show improvement in their predictive ability. Prediction at test time with text or graph representations alone in the joint model is better than those

<sup>5</sup>source code used in this paper will be published

Model	Split	Graph	User	G+U	Text	G+T	G+U+T
GCN + SkipThought	Rand	89.95	81.49	89.75	70.61	90.34	91.02
	Event	89.40	79.06	89.64	69.16	90.15	90.78
	Time	84.95	76.59	85.30	64.12	84.09	86.25
	Dist	67.78	45.30	70.03	58.68	69.82	70.66
GCN + HLSTM	Rand	89.03	83.66	88.57	86.84	91.48	<b>91.74</b>
	Event	89.34	80.22	88.62	88.39	91.69	<b>91.72</b>
	Time	84.83	74.50	85.09	81.36	85.57	<b>86.21</b>
	Dist	71.74	69.39	71.16	61.13	<b>72.16</b>	71.85

Table 4: Results of Joint Model Combining Text and Graph Relations

Model	Split	Graph	User	G+U	Text	G+T	G+U+T
GCN + HLSTM (50%)	Rand	86.73	78.62	86.24	85.62	89.31	89.35
	Event	86.55	78.34	85.89	84.52	89.21	89.51
	Time	82.25	70.93	81.45	80.05	85.57	85.48
GCN + HLSTM (10%)	Rand	76.13	57.76	75.55	78.61	81.35	81.49
	Event	76.58	57.10	75.75	77.60	80.55	80.93
	Time	73.24	54.09	72.48	72.92	76.52	76.75

Table 5: Results of Joint Model with Reduced Links for Test Documents

Text	Joint	Gold	Title
Right	Right	Right	Hacked Powell email reveals Hillary hates Obama for 2008
Right	Right	Right	Donald Trump will let James Comey testify
Center	Center	Center	Clinton: I am done with being a candidate
Center	Center	Center	Senate confirms Sessions as attorney general
Left	Left	Left	Clinton: Trump Doesn't See President Obama as an American Video
Left	Left	Left	Trump uses Twitter to promote leaked intelligence on North Korea
Center	Left	Left	Hillary Clinton's Campaign Says It Will Participate In Wisconsin Recount
Left	Center	Center	Supreme Court justices hint at striking Voting Rights Act provision
Left	Center	Right	Boston Marathon bombs: how investigators use technology to identify suspects
Right	Right	Left	Israel risks becoming apartheid state if peace talks fail, says John Kerry

Table 6: Examples of Bias Prediction by Text and Joint Model

listed in Table 2 and 3, especially for text. Note that the increase in accuracy is much greater when HLSTM model is utilized to represent text content since it is much more expressive. Making prediction with the aggregation of multiple scores usually leads to better accuracy. Another thing worth noting is that the performance of distant supervised setting is almost comparable with text classification results in full supervised setting, indicating the strength of our joint model.

To show our model's ability to handle news articles with different amounts of social information available. We train and test our joint model with only 50% and 10% of the links for test articles kept. The results are summarized in Table 5. Clearly the performance improves as more social links are available. However, even with little social links provided in the latter case, our joint model propagates information effectively and results in great increase in performance.

**Qualitative Analysis** In Table 6, we compared the bias prediction by our text and joint model on several news articles (only titles shown in the ta-

ble). These examples demonstrate how subtle the bias may be expressed in text sometimes. With the supporting information from an article's social context, our joint model can achieve more reliable prediction.

## 7 Summary

In this paper we follow the intuition that political perspectives expressed in news articles text will also be reflected in the way the documents spread and the identity of the users who endorse it. We suggest a GCN-based model capturing this social information, and show that it provides a distant supervision signal, resulting in a model performing comparably to supervised text classification models. We also study this approach in the supervised setting and show that it can significantly enhance a text-only classification model.

Modeling the broader context in which text is consumed is a vital step towards getting a better understanding of its perspective. We intend to study fine-grained political perspectives, capturing how different events are framed.



## References

- Eric Baumer, Elisha Elovic, Ying Qin, Francesca Polletta, and Geri Gay. 2015. Testing and comparing computational approaches for identifying the language of framing in political news. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1472–1482.
- Ceren Budak, Sharad Goel, and Justin M Rao. 2016. Fair and balanced? quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly*, 80(S1):250–271.
- Dallas Card, Justin Gross, Amber Boydston, and Noah A Smith. 2016. Analyzing framing through the casts of characters in the news. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1410–1420.
- Y. Choi and J. Wiebe. 2014. +/-effectwordnet: Sense-level lexicon acquisition for opinion inference. In *EMNLP*.
- David Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P. Adams. 2015. [Convolutional networks on graphs for learning molecular fingerprints](#). In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NIPS'15*, pages 2224–2232, Cambridge, MA, USA. MIT Press.
- Heba Elfardy, Mona Diab, and Chris Callison-Burch. 2015. Ideological perspective detection using semantic features. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics*, pages 137–146.
- Anjalie Field, Doron Kliger, Shuly Wintner, Jennifer Pan, Dan Jurafsky, and Yulia Tsvetkov. 2018. Framing and agenda-setting in russian news: a computational analysis of intricate political strategies. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3570–3580.
- Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. 2017. [Neural message passing for quantum chemistry](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1263–1272, International Convention Centre, Sydney, Australia. PMLR.
- Maria Glenski, Tim Weninger, and Svitlana Volkova. 2018. Identifying and understanding user reactions to deceptive and trusted social news sources. *arXiv preprint arXiv:1805.12032*.
- S. Greene and P. Resnik. 2009. More than words: Syntactic packaging and implicit sentiment. In *NAACL*.
- Aditya Grover and Jure Leskovec. 2016. [node2vec: Scalable feature learning for networks](#). In *KDD*.
- Benjamin D. Horne, William Dron, Sara Khedr, and Sibel Adali. 2018a. [Assessing the news landscape: A multi-module toolkit for evaluating the credibility of news](#). In *Companion Proceedings of the The Web Conference 2018, WWW '18*, pages 235–238, Republic and Canton of Geneva, Switzerland. International World Wide Web Conferences Steering Committee.
- Benjamin D. Horne, William Dron, Sara Khedr, and Sibel Adali. 2018b. Sampling the news producers: A large news and feature data set for the study of the complex media landscape. In *ICWSM*.
- Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014a. Political ideology detection using recursive neural networks. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1113–1122.
- Mohit Iyyer, Peter Enns, Jordan L. Boyd-Graber, and Philip Resnik. 2014b. Political ideology detection using recursive neural networks. In *ACL*.
- Kristen Johnson and Dan Goldwasser. 2016. “all i know about politics is what i read in twitter”: Weakly supervised models for extracting politicians stances from twitter. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2966–2977.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Thomas N Kipf and Max Welling. 2016. Variational graph auto-encoders. *NIPS Workshop on Bayesian Deep Learning*.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Skip-thought vectors](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3294–3302.
- Beata Beigman Klebanov, Eyal Beigman, and Daniel Diermeier. 2010. Vocabulary choice as an indicator of perspective. In *Proceedings of the ACL 2010 Conference Short Papers*, pages 253–257. Association for Computational Linguistics.

- W.H Lin, T. Wilson, J. Wiebe, and A. Hauptmann. 2006. Which side are you on?: identifying perspectives at the document and sentence levels. In *CoNLL*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13*, pages 3111–3119, USA. Curran Associates Inc.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 1003–1011. Association for Computational Linguistics.
- Fred Morstatter, Liang Wu, Uraz Yavanoglu, Stephen R Corman, and Huan Liu. 2018. Identifying framing bias in online news. *ACM Transactions on Social Computing*, 1(2):5.
- Shirui Pan, Jia Wu, Xingquan Zhu, Chengqi Zhang, and Yang Wang. 2016. Tri-party deep network representation.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch.
- A. Patwari, D. Goldwasser, and S. Bagchi. 2017. A multi-classifier system for detecting check-worthy statements in political debates. In *CIKM*.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543. ACL.
- Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *KDD*.
- Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013a. Linguistic models for analyzing and detecting biased language. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1650–1659.
- Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013b. [Linguistic models for analyzing and detecting biased language](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1650–1659. Association for Computational Linguistics.
- Filipe N Ribeiro, Lucas Henriqueo, Fabricio Benvenutoo, Abhijnan Chakraborty, Juhi Kulshrestha, Mahmoudreza Babaei, and Krishna P Gummadi. 2018. Media bias monitor: Quantifying biases of social media news outlets at large-scale.
- Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pages 593–607. Springer.
- Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015a. Line: Large-scale information network embedding. In *WWW*.
- Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015b. [Line: Large-scale information network embedding](#). In *Proceedings of the 24th International Conference on World Wide Web, WWW '15*, pages 1067–1077, Republic and Canton of Geneva, Switzerland. International World Wide Web Conferences Steering Committee.
- Svitlana Volkova, Kyle Shaffer, Jin Yea Jang, and Nathan Hodas. 2017. Separating facts from fiction: Linguistic models to classify suspicious and trusted news posts on twitter. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 647–653.
- Robert West, Hristo S Paskov, Jure Leskovec, and Christopher Potts. 2014. Exploiting social network structure for person-to-person sentiment analysis. *Transactions of the Association for Computational Linguistics*, 2:297–310.
- Yi Yang, Ming-Wei Chang, and Jacob Eisenstein. 2016a. Toward socially-infused information extraction: Embedding authors, mentions, and entities. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1452–1461.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J. Smola, and Eduard H. Hovy. 2016b. Hierarchical attention networks for document classification. In *HLT-NAACL*.
- Daniel Xiaodan Zhou, Paul Resnick, and Qiaozhu Mei. 2011. Classifying the political leaning of news articles and users from user votes. In *Proceedings of ICWSM*.