

Semi-supervised Content-based Detection of Misinformation via Tensor Embeddings

Gisel Bastidas Guacho
University of California Riverside
gbast001@ucr.edu

Neil Shah
Snap Inc.
nshah@snap.com

Sara Abdali
University of California Riverside
sabda005@ucr.edu

Evangelos E. Papalexakis
University of California Riverside
epapalex@cs.ucr.edu

ABSTRACT

Fake news may be intentionally created to promote economic, political and social interests, and can lead to negative impacts on humans beliefs and decisions. Hence, detection of fake news is an emerging problem that has become extremely prevalent during the last few years. Most existing works on this topic focus on manual feature extraction and supervised classification models leveraging a large number of labeled (fake or real) articles. In contrast, we focus on content-based detection of fake news articles, while assuming that we have a *small* amount of labels, made available by manual fact-checkers or automated sources. We argue this is a more realistic setting in the presence of massive amounts of content, most of which cannot be easily fact-checked. We represent collections of news articles as multi-dimensional tensors, leverage tensor decomposition to derive concise article embeddings that capture spatial/contextual information about each news article, and use those embeddings to create an article-by-article graph on which we propagate limited labels. Results on real-world datasets show that our method performs on par or better than existing models that are fully supervised, in that we achieve better detection accuracy using fewer labels. Specially, our proposed method achieves 75.43% of accuracy using only 30% of labels of a public dataset while a SVM-based classifier achieved 67.43%. Furthermore, our method achieves 70.92% of accuracy in a large dataset using only 2% of labels.

KEYWORDS

Fake news, tensor decomposition, semi-supervised learning, belief propagation.

1 INTRODUCTION

Misinformation on the web is a problem that has been greatly amplified by the use of social media, and the problem of fake news in particular has become ever more prevalent during the last years. Social media is a common platform for consuming and sharing news, due to its ease-of-use in diffusing content and promoting exposure/discussion. In fact, two-thirds of Americans reported getting some of their news from social media in 2017¹. Even though social media has become a news source for its advantages, it is especially vulnerable to the propagation of fake news mostly coming from unverified publishers and crowd-based content creators because

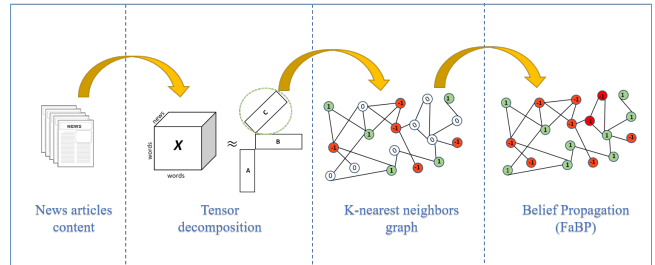


Figure 1: Our proposed method discerns real from misinformative news articles via leveraging tensor representation and semi-supervised learning in graphs.

there is practically no control over the information that is shared. The well-documented spread of misinformation on Twitter during events such as Hurricane Sandy in 2012 [4], the Boston Marathon blasts in 2013 [3] and US Presidential Elections on Facebook in 2016 [23] are all such examples. Since misinformation is intentionally created for malicious purposes such as obtain economic and political benefits or deceiving the public [22], it can clearly lead to negative user experience by either influencing their beliefs and impacting their decisions for the worse.

Several approaches in recent literature have been proposed to automatically detect misinformation using supervised classification models. Some works extract manually crafted features from news content such as the number of nouns, length of the article, fraction of positive/negative words, and more in order to discriminate fake news articles [7, 9, 19]. In addition to these works, several others proposed propagation-based models for evaluating news credibility [5, 11, 12]. Nonetheless, they initialized credibility values for the entire network using a supervised classifier. However, the reality is that such labels are often very limited and sparse. Fact-checking websites such as Snopes.com, PolitiFact.com, and FactCheck.org can be used to assess claims, but these websites require domain experts to assign credibility values to claims and are therefore, limited by human capacity. Moreover, fact-checking is a time-consuming process, often requiring surveying multiple articles and sources, evaluating reputation and likelihood of the claims before coming to a decision.

In this paper, we propose a new *semi-supervised* approach for fake news article detection based on news content, which requires *limited* labels. On a high level, our approach exploits tensor representation and decomposition of news articles, careful construction

¹<http://www.journalism.org/2017/09/07/news-use-across-social-media-platforms-2017/>

of a k -nearest neighbor graph, and propagation of limited labeled article information to conduct inference on a larger set.

Our main contributions are:

- We leverage tensor-based article embeddings, which are shown to produce concise representations of articles with respect to their spatial context, in order to derive a graph representation of news articles.
- We formulate fake news detection as a semi-supervised method that propagates known labels on a graph to determine unknown labels.
- We collect a large dataset of misinformation and real news articles publicly shared on social media.
- We evaluate our method on real datasets. Experiments on two previously used datasets demonstrate that our method outperforms prior works since it requires a fewer number of known labels and achieves comparable performance.

The remainder of this paper is organized as follows. In Section 2, we describe preliminaries and definitions used throughout this paper. Section 3 defines the problem we address. In Section 4, we introduce our proposed method in detail. In Section 5, we describe the datasets, evaluation metrics and discuss the experimental results and performance of our approach. Section 6 discusses related work, and Section 7 concludes.

2 PRELIMINARIES AND NOTATIONS

In this section, we provide preliminary definitions for technical concepts that we use throughout this paper.

2.1 CP/PARAFAC Tensor Decomposition

A tensor is a multidimensional array where its dimensions are referred as modes. The most widely used tensor decomposition is Canonical Polyadic (CP) or PARAFAC decomposition [8].

CP/PARAFAC factorizes a tensor into a sum of rank-one tensors. For instance, a three-mode tensor is decomposed into a sum of outer products of three vectors:

$$\mathcal{X} \approx \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r$$

where $\mathbf{a}_r \in \mathbb{R}^I$, $\mathbf{b}_r \in \mathbb{R}^J$, $\mathbf{c}_r \in \mathbb{R}^K$ and the outer product is given by $(\mathbf{a}_r, \mathbf{b}_r, \mathbf{c}_r)(i, j, k) = \mathbf{a}_r(i)\mathbf{b}_r(j)\mathbf{c}_r(k)$ for all i, j, k [16].

The factor matrices are defined as $\mathbf{A} = [\mathbf{a}_1 \ \mathbf{a}_2 \ \dots \ \mathbf{a}_R]$, $\mathbf{B} = [\mathbf{b}_1 \ \mathbf{b}_2 \ \dots \ \mathbf{b}_R]$, and $\mathbf{C} = [\mathbf{c}_1 \ \mathbf{c}_2 \ \dots \ \mathbf{c}_R]$ where $\mathbf{A} \in \mathbb{R}^{I \times R}$, $\mathbf{B} \in \mathbb{R}^{J \times R}$, and $\mathbf{C} \in \mathbb{R}^{K \times R}$ denote the factor matrices and R is the rank of the decomposition or the number of columns in the factor matrices.

2.2 k -nearest-neighbor graph

A k -nearest-neighbor, or k -NN graph is one in which node \mathbf{p} and node \mathbf{q} are connected by an edge if node \mathbf{p} is in \mathbf{q} 's k -nearest-neighborhood or node \mathbf{q} is in \mathbf{p} 's k -nearest-neighborhood. k is a hyperparameter that determines the density of the graph – thus, a graph constructed with small k may be sparse or poorly connected.

The k -nearest-neighbors of a point in n -dimensional space are defined using a “closeness” relation where proximity is often defined in terms of a distance metric [6] such as Euclidean ℓ_2 distance. Thus, given a set of points \mathcal{P} in n -dimensional space, a k -NN graph on \mathcal{P} can be constructed by computing the ℓ_2 distance between each

pair of points and connecting each point with the k most proximal ones.

The ℓ_2 distance d between two points \mathbf{p} and \mathbf{q} in n -dimensional space is defined as:

$$d(\mathbf{p}, \mathbf{q}) = \sqrt{\sum_{i=1}^n (q_i - p_i)^2}$$

2.3 Fast Belief Propagation (FaBP)

FaBP [13] is a fast and linearized guilt-by-association method, which improves upon the basic idea of belief propagation (BP) over a graph. In our case, a belief is the label of a news article. Hence, we use FaBP as a means to propagate label likelihood over a graph, given set of known labels. The operative intuition behind FaBP and other such guilt-by-association methods is that nodes which are “close” are likely to have similar labels or belief values.

The FaBP algorithm solves the following linear system:

$$[\mathbf{I} + a\mathbf{D} - c'\mathbf{A}]b_h = \phi_h$$

where ϕ_h and b_h denote prior and final beliefs, respectively. \mathbf{A} denotes the $n \times n$ adjacency matrix of an underlying graph of n nodes, \mathbf{I} denotes the $n \times n$ identity matrix, and \mathbf{D} is a $n \times n$ diagonal matrix of degrees where $\mathbf{D}_{ii} = \sum_j \mathbf{A}_{ij}$ and $\mathbf{D}_{ij} = 0$ for $i \neq j$.

Finally, we define $a = \frac{4h_h^2}{1-4h_h^2}$ and $c' = \frac{2h_h}{(1-4h_h^2)}$ where h_h denotes the homophily factor between nodes (i.e. their “coupling strength” or association). More specifically, higher homophily means that close nodes tend to have more similar labels. The coefficient values are set as above for convergence reasons; we refer the interested reader to [13] for further discussion.

3 PROBLEM DEFINITION

We consider a misinformative, or fake, news article as one that is “intentionally and verifiably false,” following the definition used in [22]. With this definition in mind, we aim to discern fake news articles from real ones based on their content. Henceforth, by “content,” we refer to the text of the article. We reserve the investigation of other types of content (such as image and video) for future work.

Let $\mathcal{N} = \{n_1, n_2, n_3, \dots, n_M\}$ be a collection of news articles of size M where each news article is a set of words and $\mathcal{D} = \{w_1, w_2, w_3, \dots, w_I\}$ be a dictionary of words of size I . Note that articles can have varying length. Assuming that labels of some news articles are available. Let $\mathbf{l} \in \{-1, 0, 1\}$ denote a vector containing the partially known labels, such that entries of 1 represent real articles, -1 represents fake articles and 0 denotes an unknown status. Using the aforementioned notations, we formally define the problem as follows:

Given a collection of news articles \mathcal{N} and a label vector \mathbf{l} with entries for labeled real/fake and unknown articles, we aim to **predict** the class labels of the unknown articles.

We address the problem as a binary classification problem; hence, a news article is classified either fake or real.

4 PROPOSED METHOD

In this section, we introduce a content-based method for semi-supervised fake news article detection. This method consists of three steps: representing/decomposing content as a tensor, constructing a k -NN graph of proximal embeddings, and propagating beliefs using FaBP. Figure 1 summarizes our proposed method. Below, we describe the individual steps in further detail.

Step 1: Tensor Decomposition.

Given a collection of news articles $\mathcal{N} = \{n_1, n_2, n_3, \dots, n_M\}$ of size M , where each news article in \mathcal{N} is a vector that contains the words within the news article, we build similar tensor-based article embeddings as proposed in [10]. Specifically, we propose the use of *binary-based tensor construction* method. That is, we build a three-mode tensor $\mathcal{X} \in \mathbb{R}^{I \times I \times M}$ (*words, words, news*) where for each news article, we create a co-occurrence matrix where all co-occurrence entries are boolean and indicate ($word_1, word_2$) appeared within a window parameter of w (usually 5-10) words² at least once. We then use CP/PARAFAC tensor decomposition [8] to factorize the tensor. As [10] demonstrates, such tensor-based article embeddings captures spatial/contextual nuances of different types of news articles and result in homogeneous article groups. After decomposing the tensor, we obtain the factor matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$ whose columns correspond to different latent topics, clustering news articles and words in the latent topic space. More specifically, each row of \mathbf{C} is the representation of the corresponding article in the resulting embedding space.

Step 2: k -NN graph of news articles.

The tensor embedding we computed in Step 1 provides a compact and discriminative representation of news articles into a concise set of latent topics. Using this embedding, we construct a graphical representation of news articles. In particular, we use the factor matrix \mathbf{C} in order to construct a k -NN graph G of news articles. As we mentioned before each column in \mathbf{C} is the representation of the corresponding news article in the latent topic space; thus, by constructing a k -NN graph on \mathbf{C} , we can find similar articles in that space. To this end, we consider each row in $\mathbf{C} \in \mathbb{R}^{M \times R}$ as a point in R -dimensional space. We then compute ℓ_2 distance among news and find the k -closest points for each point in \mathbf{C} . Since the number of news articles can be extremely large in practice, we can leverage well-known kd -tree based optimizations [17] in order to more efficiently find to k -nearest-neighbors for each article.

Each node in G represents a news article and each edge encodes that two articles are similar in the embedding space. In this step, we only leverage the distance as a means to measure similarity between news articles, without much concern for the actual order of proximity. Thus, we enforce symmetry in the neighborhood relations, that is, if n_1 is a k -nearest-neighbor of news n_2 , the opposite should also hold. The resultant graph G is an undirected, symmetric graph where each node is connected to at least k nodes. The graph can be compactly represented as an $M \times M$ adjacency matrix.

²We experimented with small values of that window and results were qualitatively similar.

Step 3: Belief Propagation.

Using the graphical representation of the news articles above, and considering that for a small set of those news articles we have ground truth labels, our problem becomes an instance of semi-supervised learning over graphs. We use a belief propagation algorithm which assumes homophily, because news articles that are connected in the k -NN graph are likely to be of the same type due to the construction method of the tensor embeddings; moreover, [10] demonstrates that such embeddings produces fairly homogeneous article groups. More specifically, we use the fast and linearized FaBP variant proposed in [13]. The algorithm is demonstrated to be insensitive to the magnitude of the known labels; thus, we consider that FaBP can achieve good learning performance only using a small number of known labels. Hence, we classify fake news articles in a semi-supervised fashion.

Table 1: Dataset specifics.

Datasets	# fake news	# real news	# total
Dataset1 (Political)	75	75	150
Dataset2 (Bulgarian)	69	68	137
Our dataset	31,739	31,739	63,478

5 EXPERIMENTAL EVALUATION

In this section, we evaluate the performance of our method on real datasets. We implemented our method in MATLAB using Tensor Toolbox [1] and leveraged existing MATLAB FaBP implementation [13].

5.1 Dataset description

In order to evaluate the performance of our proposed method, we used data from three datasets: two public datasets of hundreds of articles, and our own, collected dataset of tens of thousands of articles, as shown in Table 1.

5.1.1 Public datasets. The two public datasets were used in previous studies. Specifically, *Dataset1* consists of 150 political news articles, balanced to have 75 articles of each class, and was provided by [9]. *Dataset2* contains 68 real and 69 fake news articles, and was provided by [7].

5.1.2 Our dataset. In constructing our dataset, we collected news article URLs from Twitter tweets during a 3-month period from June-August 2017. These URLs were filtered based on website domain. We then crawled those URLs to get the news article content. To that end, we used web API boilerpipe³, Python library Newspaper3k⁴, and Diffbot⁵. All real news articles were featured on 367 domains obtained from Alexa⁶, and fake news articles belong to 367 domains from the BSDetector (browser extension to identify fake news sites) domain list [2]. Table 2 describes the BSDetector-specified domain categories. In this study, we considered news from those domain categories as fake news in order to

³<http://boilerpipe-web.appspot.com/>

⁴<http://newspaper.readthedocs.io/en/latest/>

⁵<https://www.diffbot.com/dev/docs/article/>

⁶<https://www.alexa.com/>

later perform the binary classification task. All in all, the dataset consists of 31,739 fake news and 409,076 real news articles. Since this dataset is highly imbalanced, we randomly down-sampled the real class to get a balanced dataset. Our resultant dataset contains 31,739 articles of each class. Figure 2 shows distribution of articles in our dataset per domain category.

The content of each news article for the three datasets was pre-processed using tokenization and stemming. Besides, stopwords and punctuations were removed.

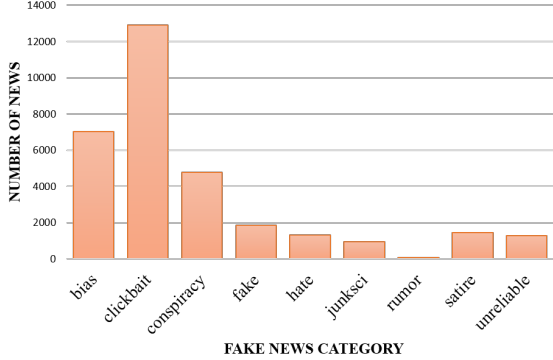


Figure 2: Distribution of misinformation per domain category in our collected dataset.

5.2 Evaluation Metrics

Since we consider the fake article detection as a binary classification problem, we evaluated our method in terms of the following commonly used metrics:

- $Accuracy = \frac{|TP|+|TN|}{|TP|+|TN|+|FP|+|FN|}$, which is the percentage of correctly classified articles.
- $Precision = \frac{|TP|}{|TP|+|FP|}$, which is the percentage of articles predicted as fake out of all articles predicted as fake.
- $Recall = \frac{|TP|}{|TP|+|FN|}$, which is the percentage of all fake articles that are correctly predicted as fake.
- $F1_score = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$, which is the harmonic mean of precision and recall, and indicates a combined measure of performance.

TP denotes true positives (i.e. correctly predicted fake news articles), FP denotes false positives (i.e. real news articles which were predicted as fake), TN denotes true negatives (i.e. correctly predicted real news articles), and FN denotes false negatives (i.e. fake news articles which were predicted as real).

5.3 Evaluation

In order to find the best-performing parameters for our method, we run an iterative process using cross-validation where we evaluated different settings with respect to R (i.e. decomposition rank) and k (i.e. the number of nearest neighbors, controlling the density of the k -NN graph). We considered values of R from 1 to 20, since decomposition rank is often set to be low for time and space reasons in practice [21]. Likewise, we tested k with values from 1 to 100,

Table 2: Domain categories collected from BSDetector [2], as indicated in our dataset.

Category	Description
<i>Bias</i>	Sources that traffic in political propaganda and gross distortions of fact.
<i>Clickbait</i>	Sources that are aimed at generating online advertising revenue and rely on sensationalist headlines or eye-catching pictures.
<i>Conspiracy</i>	Sources that are well-known promoters of kooky conspiracy theories.
<i>Fake</i>	Sources that fabricate stories out of whole cloth with the intent of pranking the public.
<i>Hate</i>	Sources that actively promote racism, misogyny, homophobia, and other forms of discrimination.
<i>Junk Science</i>	Sources that promote pseudoscience, metaphysics, naturalistic fallacies, and other scientifically dubious claims.
<i>Rumor</i>	Sources that traffic in rumors, innuendo, and unverified claims.
<i>Satire</i>	Sources that provide humorous commentary on current events in the form of fake news.

trading off greater bias for less variance with increasing k . We found that the best accuracy is obtained when both parameters R and k are set to be 10. We find that for values of k and R greater than 10, performance is qualitatively similar as shown in Figure 3, and thus we fix the parameters as such in evaluation. Notice that using a small k value (e.g. 1 or 2), the accuracy is relatively poor; this is because building a k -NN graph with small k results in a highly sparse graph which offers limited propagation capacity. In all experiments, we tested accuracy over the test set of all articles whose labels were “unknown” in the propagation step.




Figure 3: Performance using different parameter settings for decomposition rank (R) and number of nearest neighbors (k).

We evaluated our method with different percentages p of known labels. Table 3 shows the performance of our method using $p \in \{5\%, 10\%, 20\%, 30\%\}$ of labeled news articles from our dataset. Our results demonstrate that we can achieve an accuracy of 70.76% only using 10% of labeled articles. We also evaluated the performance of our approach using extremely sparse known labels. That is, we evaluated our method using $p < 5\%$ and varying the number of

nearest neighbors. Figure 4 shows that we can achieve an accuracy of 70.92% using 2% of known labels when k is set to be 200. In fact, the performance of our approach degrades fairly gracefully with even smaller proportions of known labels.

Table 3: Performance of the proposed method using our dataset with different percentages of labeled news.

%Labels	Accuracy	Precision	Recall	F1
5%	69.12 ± 0.003	69.09 ± 0.004	69.24 ± 0.009	69.16 ± 0.004
10%	70.76 ± 0.003	70.59 ± 0.003	71.13 ± 0.010	70.85 ± 0.004
20%	72.39 ± 0.001	71.95 ± 0.002	73.32 ± 0.004	72.63 ± 0.002
30%	73.44 ± 0.001	73.13 ± 0.003	74.14 ± 0.003	73.63 ± 0.001

Additionally, to evaluate the quality of tensor embeddings over traditional vectorial representations, we compared performance between our approach and a variant in which between we constructed a k -NN graph built from the term frequency inverse-document-frequency ($tf-idf$) representations. Figure 5 shows that our method with tensor embeddings consistently attains better accuracy than the alternative over varying known label percentages and suggests that binary-based tensor can better captures spatial/contextual nuances of news articles over vectorial representations.

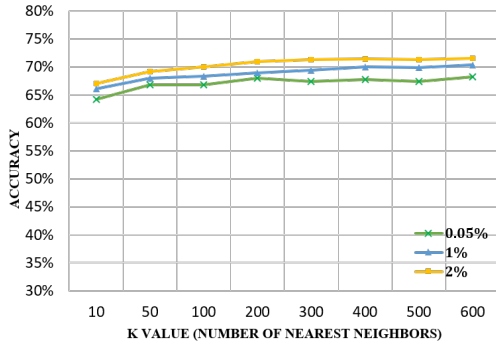


Figure 4: Performance using extremely sparse (<5%) labeled articles and varying number of nearest neighbors.

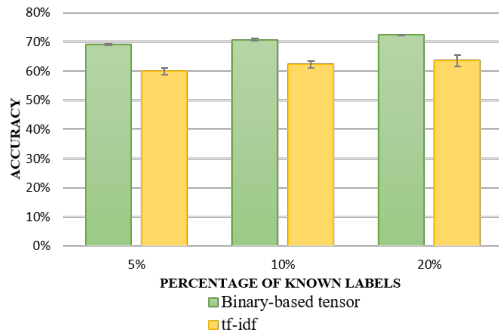


Figure 5: Performance our method using tensor-based article embedding compared to Belief propagation using a graph built from $tf-idf$ matrix.

In addition, we evaluated our model using *Dataset1* and *Dataset2*. We compare the accuracy achieved by our method to the accuracy achieved by the following approaches:

- *SVM on content-based features* as proposed in [9]. To this extent, we replicated the feature extraction from news content and used SVM in order to show the performance using different percentages of data to train the model.
- *Logistic regression on content-based features* proposed by [7]. We used their publicly available implementation. In particular, we run their method with linguistic (n -gram) feature extraction using different percentages of data to train their model.

Figure 6 shows the results for *Dataset1*. Our approach demonstrates improved accuracy even with fewer labels – specifically, we achieved 75.43% accuracy using only the 30% of news labels while SVM(30%/70% train/test), SVM(5-fold cross-validation), and logistic regression (30%/70% train/test) attained 67.43%, 71% and 50.09% of accuracy, respectively. The accuracy achieved by SVM(5-fold cross-validation) was reported by Horne et al. in [9].

For *Dataset2*, we run logistic regression and SVM, using 10%/90% train/test split. These approaches achieved an accuracy of 59.84% and 64.79%, respectively, compared to the 67.38% accuracy achieved by our approach, using the same percentage of labeled articles.

We note that our method is able to achieve this performance only having a small number of labeled news articles due to the semi-supervised nature of our approach; leveraging the k -NN graph and propagating beliefs allows us to exploit similarity between even unlabeled news articles, which supervised classification systems cannot leverage.

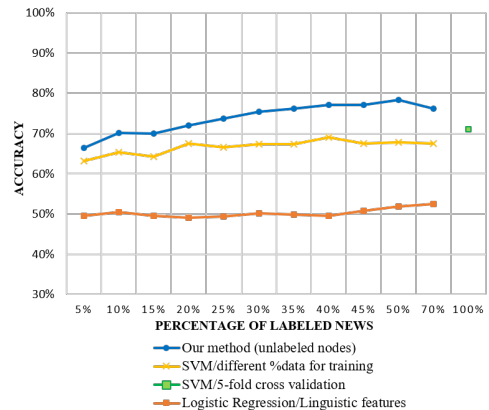


Figure 6: Performance using *Dataset1* provided by Horne et al. [9]

6 RELATED WORK

6.1 Supervised models

Several works that aim to detect misinformation are focused on supervised learning models based on extracted features from news content, user or contextual information. In [7], the authors proposed a logistic regression classifier using linguistic (n -gram), credibility (punctuation, pronoun use, capitalization) and semantic features generated from the news content. [9] used SVM on content-based features that are categorized into stylistic, complexity and

psychological features classify real, fake and satirical news. In [18], the authors propose detecting rumors by building naïve-Bayes classifiers on content, network and microblog-specific features. [14] and [20] leverage temporal structure using recurrent neural network (RNN) based models to represent text and user characteristics. In [15], the authors propose a Dynamic Series-Time Structure (DSTS) model for detecting rumors by capturing the social context of an event from content, user and propagation-based features. On the contrary, we use news content to construct a tensor of co-occurrence between words for each news article. This tensor-based representation captures similarity in the context of the spatial relations among words between news articles.

6.2 Propagation models

There are previous works that proposed propagation-based models for evaluating news credibility. [5] proposed a PageRank-like credibility propagation method on multi-typed network of events, tweets and users. [12] proposed constructing a credibility network for news verification based on conflicting viewpoints between tweets (i.e. based on positive and negative arguments about the news). In [11], the authors proposed a hierarchical propagation model on a three-layer credibility network that consists of event, sub-event and message layers. These layers form a hierarchical structure that models their relations. These works all require initial credibility values, and obtain them by using the output of a supervised classifier. Unlike these works, we focus on a semi-supervised method that leverages tensor decomposition and belief propagation, and works well with very few labels.

7 CONCLUSIONS

In this paper, we propose a semi-supervised content-based method for detecting misinformative news articles. Our method leverages tensor-based article embeddings to construct a k -nearest neighbor graph of news articles which captures similarity between them in a latent, embedding space. We then use a guilt-by-association propagation algorithm to diffuse known article labels over the graph. We evaluate our method using our own dataset of over 63K articles, and two public datasets. Experiments on these three real-world datasets demonstrate that our model is able to distinguish fake from real news only using a small number of labeled articles, compared to state-of-the-art content-based approaches which achieve similar quality while assuming fully supervised models. More specifically, our method achieves 75.43% accuracy using only 30% of labels of the first public dataset and 67.38% accuracy using only 10% of labels of the second public dataset. Additionally, our method attains 70.92% of accuracy using only 2% of labels of our dataset.

ACKNOWLEDGMENT

Research was supported by a gift from Snap Inc, an Adobe Data Science Faculty Award, and by the Department of the Navy, Naval Engineering Education Consortium under award no. N00174-17-1-0005. Any opinions, findings, and conclusions or recommendations expressed here are those of the author(s) and do not necessarily reflect the views of the funding parties.

REFERENCES

[1] T. G. Kolda B. W. Bader et al. 2015. Matlab tensor toolbox version 2.6. Available online. (February 2015).

[2] B.S. Detector. 2017. <http://bsdetecter.tech/>. (2017).

[3] A. Gupta, H. Lamba, and P. Kumaraguru. 2013. \$1.00 per RT #BostonMarathon #PrayForBoston: Analyzing Fake Content on Twitter. In *2013 APWG eCrime Researchers Summit*. 1–12. DOI: <http://dx.doi.org/10.1109/eCRS.2013.6805772>

[4] Aditi Gupta, Hemank Lamba, Ponnurangam Kumaraguru, and Anupam Joshi. 2013. Faking Sandy: Characterizing and Identifying Fake Images on Twitter During Hurricane Sandy. In *Proceedings of the 22Nd International Conference on World Wide Web (WWW '13 Companion)*. ACM, New York, NY, USA, 729–736. DOI: <http://dx.doi.org/10.1145/2487788.2488033>

[5] Manish Gupta, Peixiang Zhao, and Jiawei Han. *Evaluating Event Credibility on Twitter*. 153–164. DOI: <http://dx.doi.org/10.1137/1.9781611972825.14> arXiv:<http://epubs.siam.org/doi/pdf/10.1137/1.9781611972825.14>

[6] Jiawei Han, Micheline Kamber, and Jian Pei. 2011. *Data Mining: Concepts and Techniques* (3rd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[7] Momchil Hardalov, Ivan Koychev, and Preslav Nakov. 2016. In Search of Credible News. *Artificial Intelligence: Methodology, Systems, and Applications. AIMS 2016 Lecture Notes in Computer Science* (2016), 172fi?i180. DOI: http://dx.doi.org/10.1007/978-3-319-44748-3_17

[8] R. A. Harshman. 1970. Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis. *UCLA Working Papers in Phonetics* 16, 1 (1970), 84.

[9] Benjamin D. Horne and Sibel Adali. 2017. This Just In: Fake News Packs a Lot in Title, Uses Simpler, Repetitive Content in Text Body, More Similar to Satire than Real News. *CoRR abs/1703.09398* (2017). arXiv:1703.09398 <http://arxiv.org/abs/1703.09398>

[10] Seyedmehdi Hosseini-motlagh and Evangelos E. Papalexakis. 2017. Unsupervised Content-Based Identification of Fake News Articles with Tensor Decomposition Ensembles.

[11] Z. Jin, J. Cao, Y. G. Jiang, and Y. Zhang. 2014. News Credibility Evaluation on Microblog with a Hierarchical Propagation Model. In *2014 IEEE International Conference on Data Mining*. 230–239. DOI: <http://dx.doi.org/10.1109/ICDM.2014.91>

[12] Zhiwei Jin, Juan Cao, Yongdong Zhang, and Jiebo Luo. 2016. News Verification by Exploiting Conflicting Social Viewpoints in Microblogs. (2016). <https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12128>

[13] Danai Koutra, Tai-You Ke, U. Kang, Duen Chau, Hsing-Kuo Pao, and Christos Faloutsos. 2011. Unifying Guilt-by-Association Approaches: Theorems and Fast Algorithms. In *Machine Learning and Knowledge Discovery in Databases (ECML/PKDD)*. Lecture Notes in Computer Science, Vol. 6912. 245–260.

[14] Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard J. Jansen, Kam-Fai Wong, and Meeyoung Cha. 2016. Detecting Rumors from Microblogs with Recurrent Neural Networks. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI'16)*. AAAI Press, 3818–3824. <http://dl.acm.org/citation.cfm?id=3061053.3061153>

[15] Jing Ma, Wei Gao, Zhongyu Wei, Yueming Lu, and Kam-Fai Wong. 2015. Detect Rumors Using Time Series of Social Context Information on Microblogging Websites. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management (CIKM '15)*. ACM, New York, NY, USA, 1751–1754. DOI: <http://dx.doi.org/10.1145/2806416.2806607>

[16] Evangelos E. Papalexakis, Christos Faloutsos, and Nicholas D. Sidiropoulos. 2016. Tensors for Data Mining and Data Fusion: Models, Applications, and Scalable Algorithms. *ACM Trans. Intell. Syst. Technol.* 8, 2, Article 16 (Oct. 2016), 44 pages. DOI: <http://dx.doi.org/10.1145/2915921>

[17] Dan Pelleg and Andrew Moore. 1999. Accelerating exact k-means algorithms with geometric reasoning. In *Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 277–281.

[18] Vahed Qazvinian, Emily Rosengren, Dragomir R. Radev, and Qiaozhu Mei. 2011. Rumor Has It: Identifying Misinformation in Microblogs. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP '11)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 1589–1599. <http://dl.acm.org/citation.cfm?id=2145432.2145602>

[19] Victoria L. Rubin, Niall J. Conroy, Yimin Chen, and Sarah Cornwell. 2016. Fake News or Truth? Using Satirical Cues to Detect Potentially Misleading News.

[20] Natali Ruchansky, Sungyong Seo, and Yan Liu. 2017. CSI: A Hybrid Deep Model for Fake News. *CoRR abs/1703.06959* (2017). arXiv:1703.06959 <http://arxiv.org/abs/1703.06959>

[21] Neil Shah, Alex Beutel, Brian Gallagher, and Christos Faloutsos. 2014. Spotting suspicious link behavior with fbbox: An adversarial perspective. In *Data Mining (ICDM), 2014 IEEE International Conference on*. IEEE, 959–964.

[22] Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake News Detection on Social Media: A Data Mining Perspective. *CoRR abs/1708.01967* (2017). arXiv:1708.01967 <http://arxiv.org/abs/1708.01967>

[23] Craig Silverman. 2016. This Analysis Shows how Fake Election News Stories Outperformed Real News on Facebook. BuzzFeed News.