

Recurrent Graph Neural Networks for Rumor Detection in Online Forums

Di Huang*
dh_599@usc.edu
University of Southern California
Los Angeles, California, USA

Jacob Bartel
bartel@google.com
Google
Mountain View, California, USA

John Palowitch
palowitch@google.com
Google Research
San Francisco, California, USA

ABSTRACT

The widespread adoption of online social networks in daily life has created a pressing need for effectively classifying user-generated content. This work presents techniques for classifying linked content spread on forum websites – specifically, links to news articles or blogs – using user interaction signals alone. Importantly, online forums such as Reddit do not have a user-generated social graph, which is assumed in social network behavioral-based classification settings. Using Reddit as a case-study, we show how to obtain a derived social graph, and use this graph, Reddit post sequences, and comment trees as inputs to a Recurrent Graph Neural Network (R-GNN) encoder. We train the R-GNN on news link categorization and rumor detection, showing superior results to recent baselines. Our code is made publicly available at https://github.com/google-research/social_cascades.

KEYWORDS

graph neural networks, social networks, rumor detection, online forums

ACM Reference Format:

Di Huang, Jacob Bartel, and John Palowitch. 2021. Recurrent Graph Neural Networks for Rumor Detection in Online Forums. In *Proceedings of MIS2 workshop at KDD 2021 (KDD MIS2 '21)*. ACM, New York, NY, USA, 4 pages.

1 INTRODUCTION

As online social media becomes increasingly present in peoples' daily lives, greater proportions of users get news and journalistic content directly from their "feed" on accounts like Facebook, Twitter, YouTube, and Reddit. Pew research reported that social media outpaced print news as a news source in the United States in 2018 [11]. Some social media platforms are also news-centric – for instance, a 2016 Pew study found that seven out of ten Reddit users use the platform to get their news [1]. Following these findings, the seminal work [14] showed that user interaction signals – differential patterns of liking, re-sharing, and commenting – distinguish between posts that link to certain categories of online content, in particular content later identified as "rumors" versus other content.

*Work completed during Google Research summer internship.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

KDD MIS2 '21, August 14-18th, 2021, Virtual

© 2021 Copyright held by the owner/author(s).

Due to these phenomena, a recent sub-field of applied machine-learning research has grown, focusing on graph-based artificial intelligence models for classifying links and content shared on social media, particularly for rumor detection [3]. Most rumor detection models recently introduced have been tuned for and evaluated on data from online *social networks* like Twitter or Facebook [e.g. 9, 16]. These platforms have a natural social graph created by users, which provides an inherent graph on which a Graph Neural Network can propagate rumor information. However, relatively less attention in rumor detection research has been given to online *forums* like Reddit. Our work addresses two nuances specific to forums. First, most forums do not have a natural who-follows-who social graph. Second, most forums do not feature a "repost" option on their platform, preventing usual inter-user cascades seen in social networks [14]. Instead, each article in a forum is posted a limited number of times, each time independently by users across the platform. This means that each forum post (unlike social network posts) is the start of its own discussion cascade consisting of a long comment-tree graph, as opposed to repost/share behavior found on social networks.

To address these nuances, we provide a two-fold contribution to this space. First, we illustrate the construction of *emergent* social networks from online forum data, which we use both for feature learning and the downstream neural network computational graph via a GNN. Second, we introduce a Recurrent-GNN model which can well-handle the independent, sequential nature of article posting on forums. Our approach combines an RNN, to capture time-order of posts, with a Graph Neural Network (GNN), to capture the post comment relations of users. We evaluate our approach on the publicly-available Reddit corpus, testing two classification tasks: topic classification and rumor detection. We address preliminaries in Section 2, detail our methods in Section 3, describe our evaluation experiments in Section 4, and conclude with a brief discussion in Section 5.

2 PRELIMINARIES

News and opinion pieces are being produced with record-breaking volume, and links to these pieces propagate swiftly on social sites like Twitter, Facebook, and Reddit. Formally, a link m is propagated by a sequence of posts $\mathbf{p} = \{p_1, p_2, \dots\}$ having corresponding authors $\mathbf{a} = \{a_1, a_2, \dots\}$. On forum sites like Reddit, for which our main contributions are intended, each post also has a sequence of *commenters* $\mathbf{c} = \{c_1, c_2, \dots\}$. Note that $\mathbf{a}, \mathbf{c} \subseteq \mathbb{U}$ where \mathbb{U} is the total user set of the forum. For each task, we pair each message m with a categorical label y , and train various models to predict y given feature data associated with any \mathbf{p} , \mathbf{a} , and \mathbf{c} .

Classification tasks having the above general setting, in which a label y is inferred from a some social media information piece

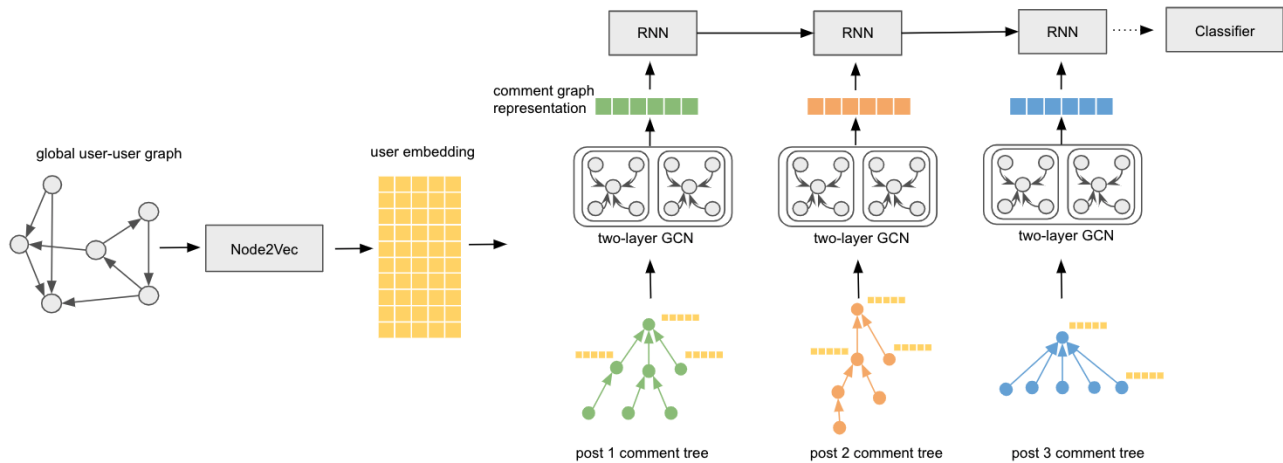


Figure 1: Recurrent Graph Neural Network (R-GNN) Framework. User features are derived from Node2Vec [5] graph embeddings. For a given article, post representations are learned from commenter features by a GCN component, which are fed to RNN cells. The RNN encodes the article representation its post representations, modelling the article’s temporal propagation.

m , have attracted great attention from researchers in recent times. For instance, [8] uses sentiment analysis to distinguish spam comments and trustworthy comments, and [12] detect fake news via assessment on credibility of article headlines. However, content-based classification in this domain can be challenging, especially in the context of rumor detection, as pieces with differing labels can nonetheless feature similar topics and writing styles. [17]. Because of this, other approaches like Traceminer [16] and CSI [10] have been proposed which learn from user interaction signals and information propagation paths. Similarly, graph embedding methods such as Node2Vec [5], SDNE [15] and graph convolutional networks – GCNs – have been widely used for network analysis and graph feature extraction in social network studies. In this work, we propose a framework which combines GCNs (Graph Convolutional Networks [e.g. 6, 7, 13]) and RNNs (Recurrent Neural Networks) to model the information diffusion process of article links.

We note that the pipeline and model we introduce for our Reddit case-study, as with many of the approaches listed above, can be used for any supervised classification tasks featuring labelled posts or article links. When applied to a task like rumor detection, it is crucial to note that a machine-learning model cannot be used to establish or predict an ultimate verdict on whether a piece of content is true, false, or of high/low quality. As detailed in Section 4, we derive a rumor label for an article by its presence on a fact-check site, regardless of verdict. In this case, our hypothesis is simply that user interaction signals in social media data can be correlated, via deep learning, to articles’ potentials to be controversial in the specific sense that they are noticed by a fact-check site. Since we train our model (and baseline models) on fact-check site data, our results are subject to any biases or errors in that data. Nevertheless, as seen in our experimental results, behavioral signals are useful in this regard, and thus our illustration of methods in this space may interest other researchers studying these phenomena.

3 METHOD

In this work we introduce a Recurrent Graph Neural Network (R-GNN), illustrated in Figure 1, which models *two* information diffusion processes on forums. First, a graph convolutional network (GCN) encodes features from commenters under each post. Second, a recurrent neural network (RNN) learns a link representation from the sequence of post encodings. We detail these components and feature construction in the next sections.

3.1 Global User-User Interaction Graph

Unlike social network (SN) platforms like Facebook and Twitter, forums like Reddit commonly do not have a natural user-friendship graph. To replace the natural graph used in SN studies [e.g. 2, 16], we *derive* a graph \mathbb{G} from user interactions. Formally, $\mathbb{G} = (\mathbb{U}, \mathbb{E})$ where \mathbb{U} is the user set and \mathbb{E} is the edge set. \mathbb{E} is a set of undirected, weighted edges $(\{u_i, u_j\}, w_{ij})$ where w_{ij} is the count of comment-replies or post-replies between u_i and u_j , on *any* post from Reddit (including those not contained in our link dataset). This graph represents *proximal* friendships between users given their commenting activity. We encode these friendships as feature inputs to our model by computing user graph embeddings on \mathbb{G} with node2vec [5].

3.2 GCN Post Encoding

In addition to the global graph \mathbb{G} , we also construct a *local reply-graph* G_p for each post p . The graph $G_p = (U_p, E_p)$ consists of the users U_p who commented on p (including the author), and each weighted edge represents the number of times each pair of commenters replied to each other. With the aforementioned graph embeddings as user features, we encode each post p into a hidden vector \mathbf{v}_p with a two-layer GCN [7] applied to G_p .

3.3 RNN+GCN Post-Sequence Encoding

We formulate inference on a link m as a temporal sequence classification problem on its time-ordered posts $\{p_{m1}, p_{m2}, \dots\}$. At each timestep, we encode the post with the comment-graph GCN, and pass that representation to an RNN unit. Finally, we predict y with a multi-layer perceptron (MLP) applied to the RNN encoding, as illustrated in Figure 1. Formally, given a link m and its corresponding post sequence p_1, p_2, \dots , we apply the GCN to obtain post encodings v_1, v_2, \dots , and infer a predicted \hat{y} as

$$\hat{y} = \arg \max \text{MLP}(\text{RNN}(v_1, v_2, \dots)) \quad (1)$$

4 EXPERIMENTS

In this section we describe the evaluation of our R-GNN model against five baselines on two tasks: article categorization and rumor detection. Four of our baselines are standard (non-neural) machine learning methods applied to simplified features. Our fifth baseline is an established RNN-based method called TraceMiner, which has been previously evaluated on similar tasks using Twitter data [16]. We evaluate two versions of our R-GNN against these baselines. First, we remove the GCN component from our model, concatenating all authors a and commenter sequences $\{c_1, c_2, \dots\}$ associated with a post into a single sequence, which we feed to an RNN. We refer to this version as "R-GNN(-replygraph)", as it removes the influence of the comment graph signal from the learning process. This provides an ablation study of R-GNN to better evaluate the combination of GCN and RNN components in our proposed approach. Second, we evaluate the full R-GNN as described in Section 3. Finally, we note that all of our experiments implicitly test our hypothesis that a "proximal" friendship graph can be derived from user interactions as a useful signal in these tasks, as described in Section 3. All-told, we design our experiments to answer three main research questions as follows:

- RQ1** Can signals derived purely from user interactions (absent a natural social graph) be successful in classifying links that are shared in online forums?
- RQ2** Can diffusion process modeling with deep neural networks outperform standard ML models, applied to online forums?
- RQ3** Can our RNN+GCN hybrid model outperform simpler RNN-only baselines, especially for rumor detection?

4.1 Baseline Methods

Here we describe baseline models against which we compare RGNN. The hyperparameters of all models, including both variants of RGNN, were tuned on a 10% validation set and tested on a 10% test set. The tables in this section report test set metrics.

SVM/XGBoost. To address RQ1, we compare R-GNN and TraceMiner with SVM and XGboost. We apply SVM and XGBoost directly to the average embedding vectors of all users that authored or commented on any post with a given link m . This provides a "shallow" model baseline against the neural models.

Traceminer [16]. Traceminer is a RNN-based diffusion model. It directly uses the *post-author* graph embedding as the post representation for RNN input. Importantly, Traceminer does not use any commenter or comment-tree information. We label this model using "Traceminer(author)". Furthermore, R-GNN(-replygraph) can

show whether commenters' information can improve the performance compared with Traceminer baseline, which only contains authors.

4.2 Link Categorization

For the link categorization task, we match links from the UCI News Aggregator Dataset ¹ [4] to Reddit posts which embed those links. Explicitly, for each link m in the UCI News Aggregator data, y is the topic label, and p is the sequence of Reddit posts which embed m . News links in this data are divided into four news categories: business, science/technology, entertainment and health. Based on the selected 8,220 URLs and their associated posts, we construct a global user network with 77.2k nodes and 153.6k edges.

In Table 1, the best performance is notated in bold and the second best score is underlined. We see that on the categorization task, SVM with author and commenter embedding ranks first based on micro-F1 score while XGoost with author and commenter embedding achieves the best macro-F1 score. However, both R-GNN(-replygraph) and Traceminer(author) are outperformed by standard baselines, and our full R-GNN achieves only comparable results. Among the diffusion-based deep learning methods, the performance of R-GNN(-replygraph) and R-GNN both surpass the baseline Traceminer(author). In addition, models with both authors' and commenters' embedding have the better results than models with solely authors information. Thus, we can see that commenters' features are highly useful for categorizing types of news.

Table 1: URL Categorization.

Model	micro-F1	macro-F1
SVM(author)	0.5757	0.4548
XGBoost(author)	0.5697	0.4649
SVM(author+commenter)	0.5953	<u>0.4770</u>
XGoost(author+commenter)	<u>0.5903</u>	0.4834
Traceminer(author)	<u>0.5182</u>	0.4055
R-GNN(-replygraph)	0.5487	0.4423
R-GNN	0.5243	0.4768

4.3 Rumor Detection

Fact-checking websites aggregate evidence for or against particular claims made by news articles and (sometimes) social media posts. Whether or not it is true, such content can be interpreted as a rumor since it caused doubt and suspicion during its propagation process. In our rumor detection task setup, we regard all the news links which show up on Snopes ², Politifact ³ and Emergent ⁴ as rumor news. The complete dataset of links is maintained at Kaggle ⁵. To find non-rumors, we use negative sampling to extract the same amount of news links from the UCI dataset to produce the URL categorization task. All told, we build a dataset of 7,352 news links, with an equal amount of postive "rumor" examples and negative examples. The global network built from the authors

¹<https://archive.ics.uci.edu/ml/datasets/News+Aggregatorused>

²<https://www.snopes.com/>

³<https://www.politifact.com/>

⁴<http://www.emergent.info/>

⁵<https://www.kaggle.com/arminehn/rumor-citation>

and commenters on Reddit posts containing these links has 201.1k nodes and 413.0k edges.

Table 2 shows that our R-GNN model has the highest F1 score and our R-GNN(-replygraph) achieves the highest accuracy. SVM(author) and Traceminer(author) ranks the second on accuracy and F1 score respectively. Overall, sequential modeling with deep learning achieved better performance than non-neural baselines on this task.

Table 2: Rumor Detection.

Model	Accuracy	F1
SVM(author)	0.6963	0.7025
XGBoost(author)	<u>0.6908</u>	0.6886
SVM(author+commenter)	0.6790	0.6447
XGBoost(author+commenter)	0.6646	0.6594
Traceminer(author)	0.6401	<u>0.7536</u>
R-GNN(-replygraph)	0.7057	0.7485
R-GNN	0.6609	0.7731

4.4 Analysis

Returning to our three research questions, we note that the features for all models were computed solely from a user-interaction based *derived* proximal friendship graph, described in Section 3. As all models performed far better than random chance on each task, we can answer RQ1 in the affirmative, that this derived graph provides a useful signal for link classification in online forums.

On the link categorization task, interestingly, non-neural baselines outperformed TraceMiner and R-GNN. For this experiment, we can tentatively answer RQ2 in the negative. This could be because the diffusion/user-reply processes for the standard news links in UCI Aggregator data news URLs across categories may be similar, and thus non-informative to the categorization task. However, an interesting find that arose from this experiment was that methods that included *commenter* features strongly outperformed those that did not (including R-GNN models vs TraceMiner). We conjecture that is likely due to the strong forum-community signals provided by the commenters' node2vec graph embeddings.

On the other hand, for the rumor detection task, our full R-GNN model outperformed both neural and non-neural baselines, and R-GNN(-replygraph). This suggests that the diffusion/user-reply processes which feed into the GCN are more useful signals for this task. Thus in this case we can answer RQ2 and RQ3 in the affirmative.

5 DISCUSSION

In this work we introduced an approach for rumor detection and (more generally) link classification on forum websites, and evaluated this approach on Reddit data. Our model, a Recurrent Graph Neural Network (RGNN), is able to capture both the diffusion process of each link through post comment-graphs via a GCN, and simultaneously the sequential nature of link-posting on forums via an RNN. When applied to a link topic categorization task, our approach had superior performance to other RNN-based methods, but had comparable performance to non-neural baselines. When applied to a rumor detection task, our approach had superior performance to all baselines. To our knowledge, this is the first appearance

of an RGNN in this space, and among the first demonstrations of deep learning on online user interactions without a natural social graph.

Automated rumor detection via artificial intelligence (and more generally, online content categorization) in social networks is a growing area of research, featuring a rich landscape of model architectures and classification tasks. In this short paper we have examined a narrow subset of potential tasks, model architectures, and available features in this space. For instance, to better understand the effect of various *interaction*-based graph signals – which let the model learn from diffusion processes on the directed graph of user actions – we have disregarded the many *content*-based signals, e.g. text or images, available for the tasks in our paper and others. However, in doing so, we have shed light on the capacity of state-of-the-art GNNs to model article sharing in forums with interaction-based features alone. This exposes headroom to improve GNN architectures and interaction-based feature construction for classification tasks.

REFERENCES

- [1] Michael Barthel, Galen Stocking, Jesse Holcomb, and Amy Mitchell. 2016. Seven-in-Ten Reddit users get news on the site. *Pew Research Center Report*. (25 February, 2016). Retrieved May 26 (2016).
- [2] Tian Bian, Xi Xiao, Tingyang Xu, Peilin Zhao, Wenbing Huang, Yu Rong, and Junzhou Huang. 2020. Rumor detection on social media with bi-directional graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 549–556.
- [3] Alessandro Bondielli and Francesco Marcelloni. 2019. A survey on fake news and rumour detection techniques. *Information Sciences* 497 (2019), 38–55.
- [4] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- [5] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable Feature Learning for Networks. In *KDD*.
- [6] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in neural information processing systems*. 1024–1034.
- [7] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. *ICLR* (2017).
- [8] Myle Ott, Claire Cardie, and Jeffrey T Hancock. 2013. Negative deceptive opinion spam. In *Proceedings of the 2013 conference of the north american chapter of the association for computational linguistics: human language technologies*. 497–501.
- [9] Nir Rosenfeld, Aron Szanto, and David C Parkes. 2020. A Kernel of Truth: Determining Rumor Veracity on Twitter by Diffusion Pattern Alone. In *Proceedings of The Web Conference 2020*. 1018–1028.
- [10] Natali Ruchansky, Sungyong Seo, and Yan Liu. 2017. Csi: A hybrid deep model for fake news detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 797–806.
- [11] Elisa Shearer. 2018. Social media outpaces print newspapers in the US as a news source. *Pew research center* 10 (2018).
- [12] Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD explorations newsletter* 19, 1 (2017), 22–36.
- [13] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).
- [14] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151.
- [15] Daixin Wang, Peng Cui, and Wenwu Zhu. 2016. Structural deep network embedding. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 1225–1234.
- [16] Liang Wu and Huan Liu. 2018. Tracing fake-news footprints: Characterizing social media messages by how they propagate. In *Proceedings of the eleventh ACM international conference on Web Search and Data Mining*. 637–645.
- [17] Zhixuan Zhou, Huankang Guan, Meghana Moorthy Bhat, and Justin Hsu. 2019. Fake news detection via NLP is vulnerable to adversarial attacks. *arXiv preprint arXiv:1901.09657* (2019).