

IRUGCN: A Graph Convolutional Network Rumor Detection Model Incorporating User Behavior*

Shu Zhou^{1,†}, Hao Wang^{1,*}, Zhengda Zhou^{1,†}, Haohan Yi^{1,†}, Bin Shi^{1,†}

¹ Nanjing University, No.163 Xianlin Avenue, Qixia District, Nanjing, Jiangsu, 210023, China

Abstract

This paper introduces a novel rumor detection model for social media that enhances identification accuracy by incorporating user behavior alongside traditional user features. Utilizing graph convolutional networks for user representation, a recurrent neural network for analyzing propagation tree structures, and an integrator for merging these analyses, the model adeptly captures both user behaviors and the dynamics of rumor spread. Tested on Twitter15 and Twitter16 datasets, it achieved superior accuracy rates of 85.2% and 87.3%, respectively, outperforming existing models. Although the model currently does not differentiate interaction stances between users through weighted graph edges, its integration of user behavior marks a significant advancement in precise rumor detection.

Keywords

User Behaviour, Graph Convolutional Network, Rumor Detection

1. Introduction

The proliferation of rumors on social media, highlighted by their significant impact during events such as the 2016 U.S. presidential election [1], presents challenges in terms of their identification and containment due to the necessity of extensive human resources and the potential for inaccuracies. Addressing this, current research largely focuses on deep learning-based rumor detection, emphasizing content and user attributes, yet often overlooks the critical aspect of user behavior patterns [2]. This study introduces the IRUGCN model, leveraging graph convolutional networks to analyze user behavior alongside traditional metrics for more effective rumor detection. By integrating user behavior analysis with content and propagation dynamics through a sophisticated model comprising a user encoder, a propagation tree encoder, and an integrator, the IRUGCN demonstrates superior performance on benchmark datasets like Twitter15 and Twitter16, offering a promising approach to mitigating the spread of misinformation on social media.

2. Methodology

2.1. Mission objective

The methodology centers around enhancing rumor detection in social media through the construction of a dataset composed of tuples representing declarations (tweets) and their associated users, arranged to form propagation trees and a user co-occurrence graph. The aim is to classify these tuples into categories of rumors (non-rumor, false rumor, true rumor, unconfirmed rumor) using a novel detection model.

2.2. Overall structure

2.2.1. User Encoder

Utilizes graph convolutional networks (GCN) to encode user behavior and static characteristics into a higher-order user representation. This involves processing an undirected graph comprising users linked based on their interactions, with adjacency matrix adjustments reflecting the significance of these interactions. The user features are represented as a

5th Extraction and Evaluation of Knowledge Entities from Scientific Documents (EKEE2024), April 22-26, 2024, Changchun, China and Online

* Corresponding author.

† These authors contributed equally.

✉ shuzhou@smail.nju.edu.cn (S. Zhou); ywhaowang@nju.edu.cn (H. Wang); zzd20010619@163.com (Z. D. Zhou);

yi_haohan@163.com (H. H. Yi); binshi@smail.nju.edu.cn (B. Shi)

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

matrix, and the GCN updates node feature matrices to integrate information from neighbor nodes.:

2.2.2. Propagation tree structure encoder

Employs bottom-up and top-down recurrent neural network encoders to capture the structural and semantic features of rumor propagation trees. The bottom-up approach aggregates child node representations to compute a parent node's representation, capturing long-distance interaction dependencies. Conversely, the top-down approach considers the current node's features and its parent node representation. The aim is to encode the propagation tree's structure and semantics into a vector representation.

Integrator: Combines the output of the user encoder and propagation tree encoder. It fuses the user representation with the propagation tree representation through a fully connected layer, aiming to accurately predict the category of each information statement by considering both user and propagation tree information.

2.2.3. Model training

Focuses on minimizing the cross-entropy loss between the model's predicted probability distribution and the true labels, incorporating a regularization term to balance the loss and prevent overfitting.

3. Experimentation and analysis

3.1. Experimental setup

3.1.1. Data sets and evaluation indicators

The study conducts experiments using Twitter datasets (Twitter15 and Twitter16), comprising 1381 and 1181 propagation trees respectively, and

involving hundreds of thousands of users. These datasets are categorized into non-rumor, false rumor, true rumor, and unknown rumor, and split in a 9:0.5:0.5 ratio for training, validation, and test sets. Model performance is evaluated based on overall accuracy and class-specific F₁ scores.

Implemented in Pytorch, the model employs recurrent neural networks and graph convolutional networks, with evaluation through 5-fold cross validation. Key configuration parameters include a 256-dimensional word vector, 256-dimensional hidden layers for user statistical features, behavioral information, and the integrator module, and a 256-sized batch with a 0.005 learning rate using the Adam optimizer. Experiments leverage Python 3.7 on a system with an NVIDIA Geforce RTX 2080 GPU. User age data undergoes preprocessing to remove unrealistic values.

3.2. Experimental results and analysis

3.2.1. Comparative analysis of methods

The experimental analysis validates the IRUGCN model's performance in rumor detection on Twitter datasets, comparing it with existing methods like BERT[3], Transformer[4], RvNN[5], UMLARD[6], and DDGCN[7]. IRUGCN, which incorporates both user behavior and structural propagation features, significantly outperforms these methods, with the top-down encoder variant showing superior results. This highlights the importance of integrating user behavioral data for accurate rumor identification.

Comparative tests reveal that direct use of user statistical features or fully connected layers for user behavior analysis is less effective than the proposed user encoder, especially on larger datasets, underscoring the encoder's efficiency in capturing complex user interactions.

Table 1 Rumor Detection Performance Comparison Results

Twitter15 Dataset						
Model	Accuracy	Non-rumor	Fake rumor	F_1	True Rumor	Unknown
BERT	0.641	0.684	0.634		0.688	0.571
Transformer	0.708	0.695	0.728		0.759	0.653
RvNN	0.723	0.682	0.758		0.821	0.654
UMLARD	0.742	0.693	0.765		0.831	0.661
DDGCN	0.812	0.793	0.773		0.851	0.741
BU-IRUGCN	0.838	0.896	0.813		0.873	0.773
TD-IRUGCN	0.852	0.799	0.873		0.931	0.809
Twitter16 Dataset						
Model	Accuracy	Non-rumor	Fake rumor	F_1	Rumor	Unknown
BERT	0.633	0.617	0.715		0.577	0.527
Transformer	0.718	0.723	0.712		0.799	0.659
RvNN	0.737	0.662	0.743		0.835	0.708

UMLARD	0.783	0.734	0.806	0.872	0.702
DDGCN	0.817	0.798	0.778	0.856	0.746
BU-IRUGCN	0.835	0.806	0.835	0.960	0.856
TD-IRUGCN	0.873	0.816	0.856	0.970	0.856

From the results in Table 2, we can see that our model outperforms other models on these 2 datasets, especially TD-IRUGCN is more superior than BU-IRUGCN.

Compared with DDGCN and UMLARD methods, the model in this paper considers not only the statistical characteristics of users but also their behavioral information. Since the group behavior of users is more likely to spread rumors, the accuracy of rumor detection can be significantly improved when user behavioral information is taken into account, although it may also bring about other rumor detection interferences. In contrast, BERT and

Transformer have relatively low accuracy because they do not integrate many user features and other propagation features.

3.2.2. Encoder Effectiveness Analysis

Further analysis of the user encoder's effectiveness highlighted its superiority over benchmark methods relying solely on user's statistical features or using fully connected layers for feature integration. The study revealed that the user encoder performs exceptionally well on larger datasets, benefiting from rich user behavior and interaction patterns.

Table 2 Performance of bottom-up propagation tree encoder combining different features on Twitter15 dataset

Twitter15					
Model	Accuracy	Non-rumor	Fake rumor	True Rumor	Unknown
BU-Features	0.760	0.750	0.766	0.816	0.700
BU-SVD	0.775	0.763	0.773	0.859	0.707
BU-IRUGCN	0.838	0.896	0.813	0.873	0.763
Twitter16					
Model	Accuracy	Non-rumor	Fake rumor	True Rumor	Unknown
BU-Features	0.763	0.665	0.805	0.866	0.687
BU-SVD	0.786	0.710	0.829	0.892	0.682
BU-IRUGCN	0.835	0.806	0.835	0.960	0.736

Table 3 Performance of top-down propagation tree encoder combining different features on Twitter15 dataset

Twitter15					
Model	Accuracy	Non-rumor	Fake rumor	True Rumor	Unknown
TD-Features	0.889	0.832	0.917	0.972	0.840
TD-SVD	0.885	0.843	0.905	0.962	0.834
TD-IRUGCN	0.852	0.799	0.873	0.931	0.809
Twitter16					
Model	Accuracy	Non-rumor	Fake rumor	True Rumor	Unknown
TD-Features	0.904	0.838	0.881	0.924	0.882
TD-SVD	0.893	0.803	0.886	0.927	0.859
TD-IRUGCN	0.883	0.816	0.856	0.870	0.856

3.2.3. Analysis of ablation experiments

A series of ablation experiments were conducted to dissect the contribution of various model components. These experiments underscored the significance of user behavioral features and revealed a comparative advantage of the top-down propagation tree structure

encoder over the bottom-up approach, attributing this to its early incorporation of global information.

3.2.4. Early Rumor Analysis

An early rumor detection analysis further confirmed IRUGCN's effectiveness, with the model outperforming counterparts at critical early detection

time points, showcasing its potential in curbing rumor spread proactively.




Figure 1: Early rumor model performance on the Twitter15 dataset




Figure 2: Early rumor model performance on the Twitter16 dataset

3.2.5. Sample Tree Cases of False Rumor Spreading

The study also includes a qualitative analysis of false rumor propagation patterns through sample tree cases, illustrating the dynamic interplay of support, rebuttal, and skepticism in rumor spread. This underlines the model's ability to capture complex propagation dynamics, contrasting with traditional models' limitations in grasping the depth of user interactions and responses.

4. Summarize

This chapter outlines the development and validation of an innovative graph neural network-based rumor detection model that significantly enhances accuracy and real-time performance by incorporating user behavior. The model integrates three pivotal components: a user encoder, a propagation tree structure encoder, and an integrator, facilitating a multi-dimensional analysis of rumors through content, user, and propagation studies.

Empirical tests on real-world datasets have underscored the model's superiority in rumor detection accuracy over existing methodologies, laying a robust groundwork for future explorations. Moving forward, the research will delve into a more

nuanced examination of inter-user interactions, such as assigning weights to user graph edges based on user interaction stances, to refine the model's sensitivity towards intricate user relationships. Additionally, the potential of transforming datasets into graph structures is being considered to further elevate the model's performance and versatility.

Acknowledgements

This paper is supported by the National Natural Science Foundation of China under contract No. 72074108, Special Project of Nanjing University Liberal Arts Youth Interdisciplinary Team (010814370113), Jiangsu Young Social Science Talents, and Tang Scholar of Nanjing University.

References

- [1] Rahim A. Rumor Identification on Twitter Data for 2020 US Presidential Elections with BERT Model[J]. UMT Artificial Intelligence Review, 2021, 1(1): 1-1.
- [2] Gumaee A, Al-Rakhami M S, Hassan M M, et al. An Effective Approach for Rumor Detection of Arabic Tweets Using eXtreme Gradient Boosting Method[J]. ACM Transactions on Asian and Low-Resource Language Information Processing, 2022, 21(1): 1-16.
- [3] Devlin J, Chang M W, Lee K, et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding[M]. arXiv, 2018.
- [4] Vaswani A, Shazeer N, Parmar N, et al. Attention Is All You Need[M]. arXiv, 2017.
- [5] Ma J, Gao W, Wong K F. Rumor Detection on Twitter with Tree-structured Recursive Neural Networks[C]//Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Melbourne, Australia: Association for Computational Linguistics, 2018: 1980-1989.
- [6] Chen X, Zhou F, Trajcevski G, et al. Multi-view learning with distinguishable feature fusion for rumor detection[J]. Knowledge-Based Systems, 2022, 240: 108085.
- [7] Sun M, Zhang X, Zheng J, et al. DDGCN: Dual Dynamic Graph Convolutional Networks for Rumor Detection on Social Media[J]. Proceedings of the AAAI Conference on Artificial Intelligence, 2022, 36(4): 4611-4619.