

# Challenges in Pre-Training Graph Neural Networks for Context-Based Fake News Detection: An Evaluation of Current Strategies and Resource Limitations

Gregor Donabauer and Udo Kruschwitz

Information Science, University of Regensburg

Regensburg, Germany

{gregor.donabauer, udo.kruschwitz}@ur.de

## Abstract

Pre-training of neural networks has recently revolutionized the field of Natural Language Processing (NLP) and has before demonstrated its effectiveness in computer vision. At the same time, advances around the detection of fake news were mainly driven by the context-based paradigm, where different types of signals (e.g. from social media) form graph-like structures that hold contextual information apart from the news article to classify. We propose to merge these two developments by applying pre-training of Graph Neural Networks (GNNs) in the domain of context-based fake news detection. Our experiments provide an evaluation of different pre-training strategies for graph-based misinformation detection and demonstrate that transfer learning does currently not lead to significant improvements over training a model from scratch in the domain. We argue that a major current issue is the lack of suitable large-scale resources that can be used for pre-training.

**Keywords:** Fake News Detection, Social Media Networks, Graph Machine Learning

## 1. Introduction

Fake news spreading online nowadays poses one of the biggest threats to democracy (Shu, 2023), presenting a urgent requirement for the development of robust detection techniques (Zhou and Zafarani, 2020). Advances in recent years were mainly driven by context-based methods, e.g. (Wei et al., 2022; Min et al., 2022; Mehta et al., 2022). Compared to the content-based paradigm, that aims to identify fake articles based on their textual content, context-based methods include various other signals as additional information (Guo et al., 2022). This for example comprises information from social media platforms, like posts associated with news articles, sharing patterns, and the profiles of connected users which are helpful cues for the task at hand (Shu et al., 2019; Sheikh Ali et al., 2022). These different signals are often put together in graph data structures and can then be aggregated using Graph Neural Networks (GNNs), e.g. (Lu and Li, 2020; Rode-Hasinger et al., 2022). At the same time, pre-training of Large Language Models (LLMs) on massive text corpora has recently transformed the field of Natural Language Processing (NLP) as this approach equips the trained models with an understanding of general aspects of written language, e.g. (Devlin et al., 2019; Brown et al., 2020; OpenAI, 2023). These advances have led to impressive progress across the entire spectrum of NLP, also in the area of fake news detection, e.g. (Raza, 2021; Hartl and Kruschwitz, 2022). In a study conducted by Hu et al. (2020a), they presented a comprehensive evaluation of pre-training techniques within

the context of GNNs. Inspired by this progress, we propose to adopt similar strategies to learn universal attributes of heterogeneous social media context graphs relevant for the domain of context-based fake news detection.

We present a set of experiments across two fake news detection datasets evaluating how different pre-training strategies perform in the domain. Towards the end we discuss current limitations of the approach and summarize our observations and findings. In the spirit of LREC-COLING we share our implementations and pre-trained models with the community to foster reproducibility<sup>1</sup>.

## 2. Related Work

### 2.1. Graph-Based Fake News Detection

We start by providing a brief contextualization within existing related work on the context-based paradigm of graph-based fake news detection. Sometimes context refers to graph-like conversation patterns of posts related to news on social media platforms (Ma and Gao, 2020; Li et al., 2020; Wei et al., 2022). However, many approaches also include user (Chandra et al., 2020; Min et al., 2022; Rode-Hasinger et al., 2022) or publisher (Mehta and Goldwasser, 2021; Mehta et al., 2022) relationships as contextual information. Recently, it has been demonstrated that taking into account such information within heterogeneous structures can lead to further improvements (Nguyen et al., 2020; Li et al., 2022; Donabauer and Kruschwitz,

---

<sup>1</sup>[https://github.com/doGregor/pretrain\\_gnns\\_FakeNewsNet](https://github.com/doGregor/pretrain_gnns_FakeNewsNet)

2023). In conclusion, graph structures, particularly heterogeneous ones, offer high potential in modelling diverse dimensions of contextual information resulting in notable improvements in the domain of fake news detection.

## 2.2. Pre-Training of Graph Neural Networks

While pre-training neural networks has become standard in NLP, we will briefly summarize relevant work in graph machine learning below. The most comprehensive evaluation of pre-training strategies was performed by Hu et al. (2020a) in the domain of Biology and Chemistry. Generally, GNN pre-training is mainly employed in specific domains like recommender systems (Hao et al., 2021) and biomedical networks (Long et al., 2022). Some concentrate on large-scale heterogeneous graphs, using methods like contrastive learning-based node and schema-level pre-training (Jiang et al., 2021), node masking (Fang et al., 2022), or structural clustering with attention aggregation (Yang et al., 2022). Others target large homogeneous graphs with node and graph-level pre-training tasks (Lu et al., 2021), or generative approaches applicable to both homogeneous and heterogeneous large-scale graphs (Hu et al., 2020b). In this study, we aim to bridge the gap of applying GNN pre-training approaches within the domain of graph-based fake news detection.

## 3. Methodology

### 3.1. Data and Graph Setup

Existing approaches to context-based fake news detection are primarily associated with only a small number of datasets. This includes Twitter 15 and Twitter 16 (Ma et al., 2017), PHEME (Kochkina et al., 2018) and WEIBO (Ma et al., 2016), which are all very limited in terms of size and social context provided. Apart from that there is FakeNewsNet (Shu et al., 2020) which stands out as the largest dataset in the domain and offers a comparatively large amount of contextual features. This dataset comprises news articles that have been fact-checked by reputable fact-checking organizations and were found circulating on social media platforms. Furthermore, it includes identifiers for the associated social media data, allowing to retrieve and analyze this information through a recrawling processes. Recent efforts to introduce even more extensive and diverse resources were made by Barnabò et al. (2022), who provide social context cues similar to FakeNewsNet but in greater amount and even with multilingual contents. However, the vast majority of these datasets heavily depends on contextual information from access-restricted social media platforms, predominantly Twitter. Recently, new restrictions

posed on the Twitter API<sup>2</sup> are making it exceedingly challenging for researchers to gain access to these datasets at a reasonable cost or with a reasonable amount of effort. Due to these limitations, we conduct our experiments using the FakeNewsNet dataset. A useful characteristic of FakeNewsNet is that it consists of two subsets: *PolitiFact* (483 graphs) and *Gossipcop* (12,214 graphs). This allows us to conduct pre-training on one subset and perform fine-tuning on the other subset. However, the amount of data that potentially can be used for pre-training of GNNs is very limited compared to the amounts of data that are used for pre-training of LLMs in NLP. The majority of methods that incorporate social media context through graph structures, primarily target the identification of individual nodes within large graphs. These nodes hold contextual information alongside the news articles to classify. This is different to the approach of Hu et al. (2020a), who focus their methodology on a large number of small graphs, each representing an individual data point. We thus adopt our previous approach (Donabauer and Kruschwitz, 2023) to construct the social media context graphs as we formulated the problem as a graph classification task which makes the setup similar to Hu et al. (2020a). Each graph consists of three different sets of nodes: (1) the actual news article to classify at the center of the graph, (2) tweets/retweets/timeline tweets related to the article, and (3) user profiles of people linked with the aforementioned posts. Each node set is disconnected from the other sets of nodes and all initial node representations are BERT text embeddings. These node sets are linked via multiple types of edges ([*tweet, cites, article*], [*user, posts, tweet*], [*user, posts, retweet*], [*retweet, cites, tweet*] and [*user, posts, timeline tweet*]) and each edge type can be interpreted as an individual adjacency matrix that connects two distinct sets of nodes. Data setup and hyperparameters in our work are adopted from our previously provided implementations (Donabauer and Kruschwitz, 2023)<sup>3</sup>. However, one issue we need to solve is that these are heterogeneous graphs, compared to homogeneous graphs used by Hu et al. (2020a). We therefore modify the pre-training strategies to be applicable to graphs with different node/edge types.

### 3.2. Pre-Training Objectives

The pre-training objectives we adopt have been proven most effective in Hu et al. (2020a) and are twofold: we use node-level objectives and graph-level objectives. The latter can predominantly

<sup>2</sup><https://developer.twitter.com/en/docs/twitter-api>

<sup>3</sup><https://github.com/doGregor/Graph-FakeNewsNet>

be viewed as strategies aimed to learn holistic graph representations useful for the fake/real news graph classification task. On the other hand, the node-level tasks serve as regularization of the GNNs to mitigate the risk of overfitting to the self-supervised graph-level pre-training.

### 3.2.1. Node-Level Objectives

We use two different node-level pre-training objectives, namely *node masking* and *context prediction*.

**Node masking** is similar to word masking in pre-training of LLMs. We mask a proportion of 15% of all nodes with placeholder values at random. Given that the initial node attributes are continuous (BERT embeddings) we simply set all masked nodes to a uniform value of 1.0 in the dimension of the original features. Since we have different node types within our heterogeneous social media graphs, we only mask nodes of type tweet/retweet/timeline tweet which have demonstrated the biggest impact on classification performance. During pre-training we jointly train a GNN that serves as node encoder and a fully connected neural network for feature reconstruction.

For **context prediction** we first create subgraphs of the original graphs. As can be seen in Figure 1 we remove the centre node (news node) from the original graphs to create the subgraphs. We perform negative sampling with a ratio of 1 to create matching and non-matching pairs of original graphs and subgraphs. In particular, the learning objective is a binary classification task of whether context and actual graph belong to the same news node. During pre-training we jointly train one GNN that serves as node encoder of the original graph and another GNN that serves as node encoder of the context graph. Afterwards we perform mean pooling over all tweet and user nodes and predict whether the embeddings originate from the same graphs. As we are classifying the whole (sub-) graphs this is slightly different to Hu et al. (2020a).

### 3.2.2. Graph-Level Objective

As described previously, the graph-level objective serves the purpose of learning holistic graph characteristics that are useful for the upcoming fake/real news graph classification task. We compared the influence of different graph properties on the fake news detection task which revealed that the number of retweets associated with an article seems to have a comparably high influence. We therefore employ the self-supervised task of predicting the number of retweet nodes each graph holds. Since they are part of the same node set as normal tweet nodes and timeline tweet nodes, the GNN faces a challenge in determining the count of retweet nodes solely from the graph properties. Specifically, it cannot read out the number

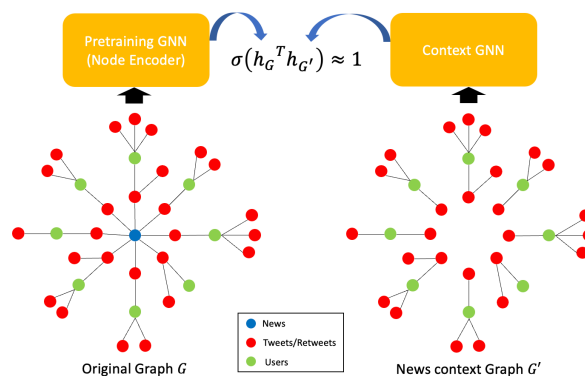


Figure 1: Formulation of pre-training on context prediction between article graph and context graph.

of retweet nodes from a dedicated node set, as these nodes are intertwined with the other tweet nodes (as described in Section 3.1). The GNN must extract this understanding by leveraging the structural information inherent in the graph, which from our point of view makes this prediction task more challenging.

## 3.3. Experiments

Inside our node encoder we use Heterogeneous Graph Transformer (HGT) GNN layers (Hu et al., 2020c) that are based on the Transformer architecture (Vaswani et al., 2017) which has proven suitable for pre-training in the domain of NLP (Devlin et al., 2019) and in other studies on pre-training of GNNs (Fang et al., 2022; Hu et al., 2020b). The encoder consists of two stacked HGT layers where each has 64 input dimensions. The GNN learns distinct sets of weights for each edge type present within the heterogeneous graph and applies them for message passing and aggregation of each node’s neighborhood to update the initial node representations. We first apply the two stacked HGT layers, and finally a fully connected layer that maps the learnt node representations to the target space (graph labels during fine-tuning or node features/context during pre-training).

During pre-training we use a batch-size of 128 and run 50 epochs for node level pre-training tasks with a learning rate of 0.001; 25 epochs (Politifact) and 50 epochs (Gossipcop) respectively for self-supervised graph-level pre-training, again with a learning rate of 0.001. For fine-tuning we again use a batch-size of 128 and a learning rate which amounts to 0.001 (same value as used during pre-training in Hu et al. (2020a)).

Our implementations are based on PyTorch Geometric (Fey and Lenssen, 2019). All experiments are executed using a single NVIDIA RTX A6000 GPU with 48 GB graphical memory size.

We evaluate how different combinations of node-level and graph-level pre-training tasks affect fake news detection performance after fine-tuning. Pre-training is always performed on one full subset (Politifact or Gossipcop), while fine-tuning is performed on 80% of the left-out subset and tested on the remaining 20% (with 5-fold cross validation). Furthermore, we are conducting experiments involving fine-tuning with limited resources. This is particularly interesting due to the high cost associated with acquiring labeled data. This additional evaluation focuses on the utilization of models pre-trained on the Gossipcop dataset as we have a greater amount of pre-training data. Subsequently, we perform fine-tuning using 50 data points (~10%) from the Politifact dataset. The model's performance is then evaluated on the complete, remaining test set, using 5-fold cross-validation.

## 4. Results

In line with common practice in NLP, we report precision, recall, accuracy and macro F1 scores for all setups (Jurafsky and Martin, 2024, Ch. 4). All values are the average results obtained by 5-fold cross validation.

As illustrated in Table 1 there are only minimal fluctuations between all the setups without any significant differences (using pairwise t-tests at  $p < 0.01$  and Bonferroni correction on the F1- and Accuracy scores derived from cross-validation across all experimental setups). The most positive impact can be seen for the pre-training strategies involving the node masking approach. For comparison, a strong state-of-the-art baseline excluding any graph-based approaches, achieves 0.956 ACC and 0.952 F1-score for Politifact as well as 0.963 ACC and 0.923 F1-score for Gossipcop (Hartl and Kruschwitz, 2022). The best-performing baseline using a graph-based approach achieves 0.975 ACC and 0.975 F1-score for Politifact as well as 0.990 ACC and 0.986 F1-score for Gossipcop (Li et al., 2022). Thus, our results on Politifact are marginally (though not significantly) better than the best-reported results in the literature.

However, the potential positive impact on model performance becomes more obvious when fine-tuning on a smaller fraction of the actual labeled data. In that case the learnt properties during pre-training lead to higher accuracy and F1 scores (compare table 2) for all setups. Again, these differences are not significant.

## 5. Discussion and Limitations

While our experiments do not demonstrate significant improvements by deploying a pre-training approach, we would like to discuss this in a broader context.

**Datasets:** The most relevant issue relates to the resources available to conduct our experiments. As already argued in Section 3.1, their sizes are notably small in comparison to the extensive amount of training data employed for pre-training in the domain of NLP. A general advantage of self-supervised pre-training is that these strategies do not require any labeled data. At the same time, online social media interactions that can serve as context are available in rich amount which means that there is a theoretical solution to this problem. However, current limitations related to the platforms' APIs are making it almost impossible to solve this problem at a reasonable cost or with a reasonable amount of effort. A potential workaround could be data augmentation by, for example, generating graphs used for pre-training at various temporal snapshots during the evolution of user interactions with a news article (Song et al., 2022). Concluding, the community working on these topics would benefit from solving the problem of limitations in access to shareable resources, similar to the practice in the area of linguistics (Cieri et al., 2022).

Another factor influencing the overall results could be that performance on FakeNewsNet is already very high, which is why choosing datasets that are more challenging and thus leave more room for improvements could better underscore the potential improvements achievable through pre-training. However, we note that our choice of FakeNewsNet for our exploratory experiments was deliberate, driven by multiple considerations: (1) its status as the largest dataset in the domain, a critical factor for effective pre-training; (2) its richness in social context cues compared to other datasets, allowing to construct heterogeneous graphs with diverse node and edge types; and (3) its division into two subsets, allowing a structured approach of pre-training on one subset followed by fine-tuning on the other subset. Examples for datasets that could be included into follow-up experiments as additional resources are the work by Mehta and Goldwasser (2023) (which is based on Baly et al. (2018) and introduces richer contextual cues) as well as the ones discussed previously such as Twitter 15 and Twitter 16 (Ma et al., 2017), PHEME (Kochkina et al., 2018), WEIBO (Ma et al., 2016) and FbMultiLingMisinfo (Barnabò et al., 2022).

**Pre-training Objectives:** As shown in Table 1, graph-level pre-training as well as context prediction do not lead to improvements while these have worked well in other domains (Hu et al., 2020a). This possibly is due to the tasks' poor alignment with the problem. For context prediction this issue becomes evident from only small fluctuations in loss and accuracy throughout the pre-training phase. One possible explanation is

Pre-training task		Fine-tuning on POL				Fine-tuning on GOS			
Node-level	Graph-level	P	R	ACC	F1	P	R	ACC	F1
–	–	0.977	0.978	0.977	0.977	0.966	<b>0.979</b>	<b>0.984</b>	0.972
ContextPred	–	0.975	0.976	0.975	0.975	0.962	0.977	0.982	0.969
NodeMasking	–	<b>0.981</b>	<b>0.982</b>	<b>0.981</b>	<b>0.981</b>	0.959	0.978	0.981	0.968
–	#RT	0.973	0.974	0.973	0.973	0.963	0.977	0.982	0.970
ContextPred	#RT	0.955	0.956	0.955	0.954	0.951	0.977	0.978	0.963
NodeMasking	#RT	0.973	0.974	0.973	0.973	<b>0.969</b>	0.978	<b>0.984</b>	<b>0.973</b>

Table 1: Precision, Recall, Accuracy and macro F1-scores for fake news detection on the *Politifact* (POL) and *Gossipcop* (GOS) datasets respectively after fine-tuning.

Pre-training task		Fine-tuning POL	
Node-level	Graph-level	ACC	F1
–	–	0.942	0.942
ContextPred	–	0.953	0.952
NodeMask.	–	<b>0.965</b>	<b>0.965</b>
–	#RT	0.957	0.957
ContextPred	#RT	0.957	0.956
NodeMask.	#RT	0.953	0.952

Table 2: Accuracy and macro F1 scores for fine-tuning with only 50 labeled samples from Politifact on models pre-trained on the Gossipcop dataset.

that, unlike in the cases of Biology and Chemistry, where node features hold high expressiveness, our situation places greater emphasis on the graph structure. This was demonstrated in our prior research, where we systematically compared the contributions of node features and graph structure to classification performance (Donabauer and Kruschwitz, 2023). However, the ego-graph structure of our social media graphs where all context converges onto a single central node is not directly comparable to molecules or proteins (which often exhibit multiple paths between nodes) and thus seems to be less fitting for this pre-training objective’s requirements. Moreover, the graph-level task evaluated might generally be less suitable for the problem of fake news detection as we could not observe any positive effects on classification performance. To address the challenges just described, one potential solution lies in employing larger, more densely connected social media context graphs. This is a common setup when framing fake news detection as a node classification task (deviating from the graph classification task alignment we maintained with Hu et al. (2020a)).

Another suggestion is to explore approaches similar to Hu et al. (2020b) who apply generative techniques to pre-train GNNs. This could even go beyond their proposed idea of learning to reconstruct the input attributed graph but move towards synthetic data generation, effectively implementing data augmentation. This direction could also

serve as a potential strategy to mitigate the problem of data scarcity described previously.

**Model Size:** While BERT’s hidden layers consist of 768 nodes (Devlin et al., 2019) and Hu et al. (2020a) used GNNs with 300 dimensional layers for pre-training, we experimented with a comparatively shallow model dimension of 64. This aligns with our goal of exploring the general effectiveness of GNN pre-training in the domain. However, using larger networks with higher capacity will potentially lead to better results.

## 6. Conclusion

We demonstrate how a current lack of large-scale resources leads to limitations in pre-training GNNs in the domain of context-based fake news detection. However, there is much potential left as we could observe improvements in experimental setups with small amounts of labelled data. Taking into account all of our observations, we proposed a range of potential avenues for future research in the field.

## 7. Ethical Considerations

We recognize the importance of responsible and ethical research practices, especially in the context of fake news detection. We provide clear documentation of our methodologies, including the pre-training strategies employed, as well as datasets and evaluation metrics used. We encourage open and transparent discussions about our findings and welcome scrutiny from the research community which is why we also share all our implementations and pre-trained models. Our research aims to contribute to the advancement of AI in a manner that benefits society, while also being aware of potential risks, challenges and limitations. We acknowledge the potential biases in data, algorithms, and models and have made efforts to mitigate these biases to the best of our knowledge. We acknowledge that fake news detection is a complex field with ethical implications. While our study focuses on technical aspects, we are aware of the broader societal impact.

## 8. Acknowledgements

We would like to thank the anonymous reviewers for their constructive feedback which has helped us improve the paper.

This work was supported by the project COURAGE: A Social Media Companion Safeguarding and Educating Students funded by the Volkswagen Foundation, grant number 95564.

## 9. References

- Ramy Baly, Georgi Karadzhov, Dimitar Alexandrov, James Glass, and Preslav Nakov. 2018. [Predicting factuality of reporting and bias of news media sources](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3528–3539, Brussels, Belgium. Association for Computational Linguistics.
- Giorgio Barnabò, Federico Siciliano, Carlos Castillo, Stefano Leonardi, Preslav Nakov, Giovanni Da San Martino, and Fabrizio Silvestri. 2022. [FbMultiLingMisinfo: Challenging Large-Scale Multilingual Benchmark for Misinformation Detection](#). In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Shantanu Chandra, Pushkar Mishra, Helen Yanakoudakis, Madhav Nimishakavi, Marzieh Saeidi, and Ekaterina Shutova. 2020. [Graph-based Modeling of Online Communities for Fake News Detection](#). *CoRR*, abs/2008.06274.
- Christopher Cieri, Mark Liberman, Sunghye Cho, Stephanie Strassel, James Fiumara, and Jonathan Wright. 2022. [Reflections on 30 Years of Language Resource Development and Sharing](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 543–550, Marseille, France. European Language Resources Association.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Gregor Donabauer and Udo Kruschwitz. 2023. [Exploring Fake News Detection with Heterogeneous Social Media Context Graphs](#). In *Advances in Information Retrieval*, pages 396–405, Cham. Springer Nature Switzerland.
- Yang Fang, Xiang Zhao, Yifan Chen, Weidong Xiao, and Maarten de Rijke. 2022. [PF-HIN: Pre-Training for Heterogeneous Information Networks](#). *IEEE Transactions on Knowledge and Data Engineering*, pages 1–14.
- Matthias Fey and Jan E. Lenssen. 2019. [Fast Graph Representation Learning with PyTorch Geometric](#). In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. [A Survey on Automated Fact-Checking](#). *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Bowen Hao, Jing Zhang, Hongzhi Yin, Cuiping Li, and Hong Chen. 2021. [Pre-Training Graph Neural Networks for Cold-Start Users and Items Representation](#). In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining, WSDM '21*, page 265–273, New York, NY, USA. Association for Computing Machinery.
- Philipp Hartl and Udo Kruschwitz. 2022. [Applying Automatic Text Summarization for Fake News Detection](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2702–2713, Marseille, France. European Language Resources Association.
- Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. 2020a. [Strategies for Pre-training Graph Neural Networks](#). In *International Conference on Learning Representations*.
- Ziniu Hu, Yuxiao Dong, Kuansan Wang, Kai-Wei Chang, and Yizhou Sun. 2020b. [GPT-GNN: Generative Pre-Training of Graph Neural Networks](#). In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20*, page 1857–1867,

- New York, NY, USA. Association for Computing Machinery.
- Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020c. [Heterogeneous Graph Transformer](#). In *Proceedings of The Web Conference 2020, WWW '20*, page 2704–2710, New York, NY, USA. Association for Computing Machinery.
- Xunqiang Jiang, Tianrui Jia, Yuan Fang, Chuan Shi, Zhe Lin, and Hui Wang. 2021. [Pre-Training on Large-Scale Heterogeneous Graph](#). In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD '21*, page 756–766, New York, NY, USA. Association for Computing Machinery.
- Daniel Jurafsky and James Martin. 2024. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. 3rd edition (draft), 3rd February 2024. <https://web.stanford.edu/~jurafsky/slp3/>.
- Elena Kochkina, Maria Liakata, and Arkaitz Zubiaga. 2018. [All-in-one: Multi-task Learning for Rumour Verification](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3402–3413, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Jiawen Li, Yudianto Sujana, and Hung-Yu Kao. 2020. [Exploiting Microblog Conversation Structures to Detect Rumors](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5420–5429, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tianle Li, Yushi Sun, Shang-ling Hsu, Yanjia Li, and Raymond Chi-Wing Wong. 2022. [Fake News Detection with Heterogeneous Transformer](#). arXiv. 10.48550/arxiv.2205.03100.
- Yahui Long, Min Wu, Yong Liu, Yuan Fang, Chee Keong Kwoh, Jinmiao Chen, Jiawei Luo, and Xiaoli Li. 2022. [Pre-training graph neural networks for link prediction in biomedical networks](#). *Bioinformatics*, 38(8):2254–2262.
- Yi-Ju Lu and Cheng-Te Li. 2020. [GCAN: Graph-aware Co-Attention Networks for Explainable Fake News Detection on Social Media](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 505–514, Online. Association for Computational Linguistics.
- Yuanfu Lu, Xunqiang Jiang, Yuan Fang, and Chuan Shi. 2021. [Learning to Pre-train Graph Neural Networks](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(5):4276–4284.
- Jing Ma and Wei Gao. 2020. [Debunking Rumors on Twitter with Tree Transformer](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5455–5466, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- 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*, page 3818–3824. AAAI Press.
- Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. [Detect Rumors in Microblog Posts Using Propagation Structure via Kernel Learning](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 708–717, Vancouver, Canada. Association for Computational Linguistics.
- Nikhil Mehta and Dan Goldwasser. 2021. [Tackling Fake News Detection by Interactively Learning Representations using Graph Neural Networks](#). In *Proceedings of the First Workshop on Interactive Learning for Natural Language Processing*, pages 46–53, Online. Association for Computational Linguistics.
- Nikhil Mehta and Dan Goldwasser. 2023. [Interactively learning social media representations improves news source factuality detection](#). In *Findings of the Association for Computational Linguistics: IJCNLP-AAACL 2023 (Findings)*, pages 291–311, Nusa Dua, Bali. Association for Computational Linguistics.
- Nikhil Mehta, Maria Pacheco, and Dan Goldwasser. 2022. [Tackling Fake News Detection by Continually Improving Social Context Representations using Graph Neural Networks](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1363–1380, Dublin, Ireland. Association for Computational Linguistics.
- Erxue Min, Yu Rong, Yatao Bian, Tingyang Xu, Peilin Zhao, Junzhou Huang, and Sophia Ananiadou. 2022. [Divide-and-Conquer: Post-User Interaction Network for Fake News Detection on Social Media](#). In *Proceedings of the ACM Web Conference 2022, WWW '22*, page 1148–1158, New York, NY, USA. Association for Computing Machinery.

- Van-Hoang Nguyen, Kazunari Sugiyama, Preslav Nakov, and Min-Yen Kan. 2020. [FANG: Leveraging Social Context for Fake News Detection Using Graph Representation](#). In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management, CIKM '20*, page 1165–1174, New York, NY, USA. Association for Computing Machinery.
- OpenAI. 2023. [GPT-4 Technical Report](#). arXiv. 10.48550/arxiv.2303.08774.
- Shaina Raza. 2021. [Automatic Fake News Detection in Political Platforms - A Transformer-based Approach](#). In *Proceedings of the 4th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2021)*, pages 68–78, Online. Association for Computational Linguistics.
- Samy Rode-Hasinger, Anna Kruspe, and Xiao Xiang Zhu. 2022. [True or False? Detecting False Information on Social Media Using Graph Neural Networks](#). In *Proceedings of the Eighth Workshop on Noisy User-generated Text (W-NUT 2022)*, pages 222–229, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Zien Sheikh Ali, Abdulaziz Al-Ali, and Tamer Elsayed. 2022. [Detecting Users Prone to Spread Fake News on Arabic Twitter](#). In *Proceedings of the 5th Workshop on Open-Source Arabic Corpora and Processing Tools with Shared Tasks on Qur'an QA and Fine-Grained Hate Speech Detection*, pages 12–22, Marseille, France. European Language Resources Association.
- Kai Shu. 2023. [Combating Disinformation on Social Media and Its Challenges: A Computational Perspective](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(13):15454–15454.
- Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2020. [FakeNewsNet: A Data Repository with News Content, Social Context, and Spatiotemporal Information for Studying Fake News on Social Media](#). *Big Data*, 8(3):171–188. PMID: 32491943.
- Kai Shu, Suhang Wang, and Huan Liu. 2019. [Beyond News Contents: The Role of Social Context for Fake News Detection](#). In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM '19*, page 312–320, New York, NY, USA. Association for Computing Machinery.
- Chenguang Song, Yiyang Teng, Yangfu Zhu, Siqi Wei, and Bin Wu. 2022. [Dynamic graph neural network for fake news detection](#). *Neurocomputing*, 505:362–374.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is All you Need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Lingwei Wei, Dou Hu, Wei Zhou, and Songlin Hu. 2022. [Uncertainty-aware Propagation Structure Reconstruction for Fake News Detection](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2759–2768, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Yaming Yang, Ziyu Guan, Zhe Wang, Wei Zhao, Cai Xu, Weigang Lu, and Jianbin Huang. 2022. [Self-supervised heterogeneous graph pre-training based on structural clustering](#). *Advances in Neural Information Processing Systems*, 35:16962–16974.
- Xinyi Zhou and Reza Zafarani. 2020. [A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities](#). *ACM Comput. Surv.*, 53(5).