International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 04 | Apr 2026 www.irjet.net p-ISSN: 2395-0072

TEMPORAL MISINFORMATION SPREAD FORECASTING IN ONLINE SOCIAL NETWORKS USING GRAPH-STRUCTURED DEEP LEARNING MODELS

Manish Kumar Pandey[1], Mrs. Arifa Khan[2]

[1]*Master of Technology, Computer Science and Engineering, Lucknow Institute of Technology, Lucknow, India*
[2]*Assistant Professor, Department of Computer Science and Engineering, Lucknow Institute of Technology, Lucknow, India*

Abstract - *The rapid growth of online social networks has significantly accelerated the spread of misinformation, posing serious challenges to public opinion, governance, and societal stability. Traditional misinformation detection approaches primarily rely on content analysis or static network structures, which are insufficient for capturing the dynamic and evolving nature of information diffusion. This study proposes a novel framework for temporal misinformation spread forecasting using graph-structured deep learning models. The approach models social networks as dynamic graphs, where users are represented as nodes and their interactions as time-dependent edges. By integrating Graph Neural Networks with temporal learning mechanisms, the proposed model effectively captures both structural relationships and temporal evolution of user interactions. The framework incorporates structural, temporal, and user-level features to learn meaningful representations of misinformation propagation patterns. Experiments are conducted on benchmark datasets, including PolitiFact, GossipCop, and ReCOVery, to evaluate the effectiveness of the proposed approach. The results demonstrate that the model outperforms traditional machine learning and static graph-based methods in terms of accuracy, precision, recall, and F1-score. Furthermore, the model shows strong capability in forecasting cascade growth and identifying early signals of misinformation spread. The findings highlight the importance of integrating temporal dynamics with graph-based learning for improving predictive performance and enabling proactive mitigation strategies in online social networks.*

Key Words: Misinformation Detection, Temporal Graphs, Graph Neural Networks, Social Network Analysis, Information Diffusion, Deep Learning

1. INTRODUCTION

The proliferation of online social networks has fundamentally transformed the way information is created, shared, and consumed. Platforms such as Twitter, Facebook, and YouTube enable real-time communication across large and diverse user communities, resulting in rapid information dissemination at an unprecedented scale. While this interconnected environment facilitates knowledge sharing and public engagement, it also introduces significant challenges related to the spread of misinformation. The increasing volume and velocity of information exchange make it difficult to monitor, verify, and control misleading content, thereby necessitating advanced computational approaches for understanding and predicting information diffusion dynamics.

1.1 Background

1.1.1 Rapid Growth of Social Networks and Information Diffusion

Online social networks have evolved into complex and large-scale systems where users interact through sharing, reposting, and commenting on content. These interactions form intricate network structures that influence how information propagates across communities. Information diffusion in such networks often follows cascade patterns, where a single piece of content spreads through multiple layers of interconnected users. The structural and temporal characteristics of these networks significantly affect the reach and speed of information dissemination (Holme and Saramäki, 2012). As a result, understanding diffusion mechanisms has become a critical research area in social network analysis and data science.

1.1.2 Emergence and Impact of Misinformation

Alongside the growth of social media, misinformation has emerged as a major global concern. False or misleading information can spread rapidly due to its often sensational nature and the absence of strict verification mechanisms on digital platforms. Studies have shown that misinformation can propagate faster and reach wider audiences compared to factual information, amplifying its societal impact (Vosoughi, Roy and Aral, 2018). The consequences of misinformation are far-reaching, affecting public opinion, political processes, and public health decisions. This highlights the urgent need for effective mechanisms to detect and control its spread.

1.2 Research Problem

1.2.1 Static Models Fail to Capture Temporal Dynamics

Traditional approaches for misinformation detection primarily rely on static representations of social networks or textual content analysis. These methods often fail to capture

the dynamic nature of user interactions, where relationships and information flows continuously evolve over time. Static models cannot effectively represent temporal changes in network structure or track the progression of information cascades, leading to limited analytical capabilities (Holme and Saramäki, 2012).

1.2.2 Lack of Forecasting Capability in Existing Approaches

Most existing studies focus on identifying misinformation after it has already spread, rather than predicting its future propagation. This reactive approach limits the ability to implement timely interventions. Forecasting misinformation spread requires models that can learn from historical interaction patterns and predict future cascade behavior, which remains a significant challenge in current research (Cheng et al., 2014).

1.3 Motivation

1.3.1 Need for Early Detection of Misinformation Cascades

Early detection of misinformation is essential to minimize its impact on society. Once misinformation reaches a large audience, it becomes increasingly difficult to correct or contain. Predictive models that can identify early signals of cascade growth enable proactive intervention strategies, such as content moderation and fact-checking, thereby reducing the spread of misleading information (Vosoughi, Roy and Aral, 2018).

1.3.2 Importance of Community-Driven Propagation

Misinformation often spreads within tightly connected communities where users share similar beliefs and reinforce each other's viewpoints. These community structures play a crucial role in amplifying and sustaining misinformation cascades. Understanding how such communities influence diffusion patterns is therefore critical for designing effective detection and forecasting models (Cinelli et al., 2021).

1.4 Research Contributions

This study makes the following key contributions:

- Proposes a temporal graph-based deep learning framework for modeling misinformation diffusion in online social networks.
- Develops a method to model dynamic interaction networks using time-evolving graph structures.
- Introduces a cascade forecasting mechanism to predict the future spread of misinformation.
- Demonstrates significant performance improvement over baseline models, including traditional machine learning and static network approaches.

1.5 Paper Organization

The remainder of this paper is organized as follows. Section 2 reviews related work on misinformation detection, graph-based learning, and temporal modeling techniques. Section 3 presents the proposed methodology, including data processing, graph construction, and model architecture. Section 4 discusses the experimental setup, results, and performance evaluation. Finally, Section 5 concludes the paper and outlines future research directions.

2. RELATED WORK

Understanding and forecasting misinformation spread in online social networks has attracted significant research attention across multiple domains, including data mining, machine learning, and network science. Existing studies have explored various approaches ranging from content analysis to advanced graph-based learning techniques. This section reviews the key developments in misinformation detection, machine learning methodologies, graph-based learning, and temporal modeling, followed by the identification of existing research gaps.

2.1 Misinformation Detection Approache

2.1.1 Content-Based, User-Based, and Network-Based Methods

Early research on misinformation detection primarily focused on content-based methods, which analyze textual features such as linguistic patterns, sentiment, and semantic inconsistencies to identify misleading information. These approaches leverage natural language processing techniques to classify news articles or social media posts as genuine or fake. While effective in certain scenarios, content-based methods often struggle with sophisticated misinformation that mimics credible writing styles (Zhou and Zafarani, 2020).

User-based approaches extend this analysis by evaluating the credibility and behavioral patterns of individuals who share information. These methods consider features such as user activity, historical reliability, and engagement patterns to identify accounts likely to propagate misinformation. However, user behavior alone may not fully capture the complexity of information diffusion, particularly in large-scale networks (Castillo, Mendoza and Poblete, 2011).

Network-based methods address these limitations by analyzing the structural relationships among users. By representing social networks as graphs, these approaches examine interaction patterns, community structures, and propagation pathways. Network-based analysis provides valuable insights into how misinformation spreads across interconnected communities, although many such methods rely on static representations of networks (Wu and Liu, 2018).

2.2 Machine Learning & Deep Learning Methods

2.2.1 Traditional Machine Learning vs Deep Learning Limitations

Machine learning techniques have been widely applied to misinformation detection due to their ability to learn patterns from large datasets. Traditional supervised learning models, such as Support Vector Machines, Decision Trees, and Logistic Regression, rely on handcrafted features derived from textual content or user metadata. While these models are computationally efficient, their performance is often limited by the quality of feature engineering and their inability to capture complex relationships within data (Ahmed, Traore and Saad, 2017).

Deep learning approaches, including Convolutional Neural Networks and Recurrent Neural Networks, have improved detection performance by automatically learning hierarchical representations from data. These models are particularly effective in capturing semantic and contextual information from textual content. However, most deep learning architectures are designed for sequential or grid-structured data and do not inherently model relational dependencies present in social networks. As a result, they may fail to capture the influence of network topology on misinformation propagation (LeCun, Bengio and Hinton, 2015).

2.3 Graph-Based Learning

2.3.1 Graph Neural Networks for Social Networks

Graph-based learning has emerged as a powerful paradigm for analyzing relational data, particularly in the context of social networks. Graph Neural Networks (GNNs) extend deep learning to graph-structured data by enabling nodes to aggregate information from their neighbors through iterative message passing. This allows the model to learn representations that capture both local and global structural patterns within the network (Hamilton, Ying and Leskovec, 2017).

In misinformation detection, GNNs have been used to model user interactions and propagation structures, enabling the identification of influential nodes and communities involved in spreading misleading information. By incorporating both node features and network topology, GNNs provide a more comprehensive understanding of information diffusion compared to traditional approaches. However, many existing GNN-based models focus on static graphs and do not account for temporal changes in user interactions (Wu et al., 2021).

2.4 Temporal Graph Models

2.4.1 Temporal Networks and Dynamic GNNs

Information diffusion in online social networks is inherently dynamic, as user interactions evolve over time. Temporal graph models address this challenge by incorporating time-dependent information into graph representations. These models capture the evolution of network structures and allow researchers to analyze how interactions change across different time intervals (Holme and Saramäki, 2012).

Dynamic Graph Neural Networks extend traditional GNNs by integrating temporal learning mechanisms, enabling the analysis of time-evolving graphs. These models process sequences of graph snapshots or continuous interaction events to learn temporal patterns in data. In the context of misinformation forecasting, temporal GNNs provide the ability to track cascade development and predict future propagation behavior. Despite their advantages, existing approaches often focus on either structural or temporal aspects rather than combining both effectively (Kazemi et al., 2020).

2.5 Research Gap

Despite significant advancements in misinformation detection and graph-based learning, a critical research gap remains in the integration of structural and temporal modeling for forecasting purposes. Many existing approaches either focus on static network structures or analyze temporal patterns independently, resulting in incomplete representations of information diffusion processes. Furthermore, most studies emphasize detection rather than prediction, limiting their ability to provide early warnings for emerging misinformation cascades.

There is a clear need for unified frameworks that combine graph-structured learning with temporal dynamics to model evolving interaction networks and forecast future propagation patterns. Such integrated approaches have the potential to improve predictive accuracy and support proactive intervention strategies in online social networks (Rossi et al., 2020).

3. PROPOSED METHODOLOGY

This section presents the proposed methodology for forecasting misinformation spread in online social networks using graph-structured deep learning models. The framework integrates data preprocessing, graph construction, feature extraction, and temporal graph-based learning to model dynamic interaction patterns and predict future propagation behavior.

3.1 Overall Framework

3.1.1 System Pipeline

The proposed framework follows a systematic pipeline consisting of multiple stages, including data collection, preprocessing, graph construction, feature extraction, model training, and prediction. Initially, raw data is collected from benchmark misinformation datasets and preprocessed to

remove noise and inconsistencies. Subsequently, social interactions are represented as graph structures, where nodes denote users and edges represent their interactions. Feature engineering techniques are then applied to extract structural, temporal, and user-level attributes. These features are used to train a graph-structured deep learning model capable of learning complex propagation patterns. Finally, the trained model predicts future misinformation spread, including cascade growth and interaction dynamics.

3.2 Data Collection and Preprocessing

3.2.1 Datasets

The study utilizes widely recognized benchmark datasets for misinformation research, including PolitiFact, GossipCop, and ReCOVery. These datasets contain labeled news content along with associated user interactions and propagation information, making them suitable for analyzing misinformation diffusion patterns across social networks.

3.2.2 Data Preprocessing

Data preprocessing is performed to ensure data quality and consistency before model development. This process begins with noise removal, where irrelevant, duplicate, or incomplete records are eliminated. Interaction extraction is then conducted to identify user activities such as sharing, reposting, and replying. Additionally, timestamp alignment is applied to organize interaction events in chronological order, enabling accurate modeling of temporal dynamics in the network.

3.3 Graph Construction

3.3.1 Static Graph Representation

In the static graph model, the social network is represented as a fixed structure where nodes correspond to users and edges represent interactions such as shares, reposts, and replies. This representation captures the structural relationships among users and provides insights into connectivity and influence patterns within the network.

3.3.2 Temporal Graph Representation

To capture the dynamic nature of social interactions, a temporal graph model is constructed by incorporating time-stamped edges. In this representation, interactions are organized as a sequence of graph snapshots over time, reflecting the evolution of the network. Temporal graphs enable the analysis of how misinformation cascades develop and propagate across different time intervals.

3.4 Feature Engineering

3.4.1 Structural Features

Structural features describe the topological characteristics of the network. Degree centrality measures the number of connections associated with each node, indicating user influence. The clustering coefficient captures the tendency of nodes to form tightly connected groups, while community density reflects the strength of interactions within specific communities. These features help identify influential nodes and structural patterns in misinformation propagation.

3.4.2 Temporal Features

Temporal features capture the dynamic aspects of information diffusion. The cascade growth rate measures how quickly misinformation spreads across the network, while interaction frequency indicates the rate of user engagement over time. Propagation duration represents the total time taken for a cascade to evolve. These features provide insights into the temporal behavior of misinformation spread.

3.4.3 User-Level Features

User-level features focus on individual behavior within the network. The influence score quantifies a user's ability to affect others, often based on connectivity and engagement. Engagement metrics, such as likes, shares, and comments, measure the level of interaction a user generates. These features help identify key contributors to misinformation diffusion.

3.5 Graph-Structured Deep Learning Model

3.5.1 Graph Neural Network Architecture

The proposed model employs a Graph Neural Network architecture to learn representations from graph-structured data. Node embeddings are generated through a message passing mechanism, where each node aggregates information from its neighboring nodes. This process enables the model to capture both local and global structural patterns within the network. Neighborhood aggregation functions combine features from connected nodes to update node representations iteratively across multiple layers.

3.5.2 Temporal Learning Mechanism

To incorporate temporal dynamics, the model integrates a temporal learning mechanism that processes sequences of graph snapshots. Dynamic graph modeling allows the system to track changes in network structure over time. Time-aware embedding updates ensure that node representations reflect both structural relationships and temporal evolution, enabling the model to capture the progression of misinformation cascades effectively.

3.5.3 Forecasting Module

The forecasting module predicts future misinformation propagation based on learned representations. It estimates cascade size by predicting the number of users likely to be influenced. Additionally, the model forecasts future interactions and identifies potential propagation paths through which misinformation may spread. This predictive

capability supports early detection and intervention strategies.

3.6 Model Training

3.6.1 Training Strategy

The dataset is divided into training, validation, and testing sets to ensure robust model evaluation. The training set is used to learn model parameters, while the validation set is used for hyperparameter tuning. The testing set evaluates the model's generalization capability on unseen data.

3.6.2 Optimization and Implementation

The model is trained using optimization algorithms such as Stochastic Gradient Descent or Adam to minimize prediction error. Implementation is carried out using Python-based deep learning frameworks, specifically PyTorch and PyTorch Geometric, which provide efficient tools for graph-based computation and model development.

4. EXPERIMENTAL RESULTS AND DISCUSSION

This section presents the experimental evaluation of the proposed temporal graph-based deep learning model for forecasting misinformation spread. The results demonstrate the effectiveness of the model in capturing both structural and temporal dynamics of social networks and compare its performance with baseline approaches.

4.1 Experimental Setup

4.1.1 Hardware and Software Configuration

The experiments are conducted in a computational environment designed to efficiently support deep learning and graph-based processing. The system is equipped with a high-performance processor, sufficient memory, and GPU acceleration to handle large-scale social network data. The implementation is carried out using Python as the primary programming language due to its extensive ecosystem for machine learning and data analysis. Deep learning models are developed using the PyTorch framework, while PyTorch Geometric is utilized for graph-based computations. This setup ensures efficient training and scalability of the proposed model.

4.1.2 Experimental Environment Details

The configuration of the experimental environment is summarized in Table .1.

Table 1: Experimental Setup

S.No	Component	Configuration
1	Operating System	Windows 11
2	Programming Language	Python 3.x
3	Framework	PyTorch
4	Graph Library	PyTorch Geometric
5	CPU	Intel i7
6	RAM	16 GB or higher
7	GPU	NVIDIA CUDA-enabled GPU

4.2 Performance Evaluation

4.2.1 Results on Benchmark Datasets

The proposed model is evaluated using benchmark misinformation datasets, including PolitiFact, GossipCop, and ReCOVery. These datasets provide labeled information along with user interaction data, enabling comprehensive analysis of misinformation propagation. The model is assessed using classification metrics such as accuracy, precision, recall, and F1-score.

The results indicate that the proposed temporal graph-based model achieves high predictive performance across all datasets. The integration of structural and temporal features enables the model to effectively capture complex propagation patterns, resulting in improved classification and forecasting accuracy. Additionally, the model demonstrates strong capability in predicting cascade growth and identifying early signals of misinformation spread.

4.2.2 Quantitative Performance Analysis

The overall performance of the model across different datasets is presented in Table .2.

Table 2: Performance on Benchmark Datasets

Dataset	Accuracy	Precision	Recall	F1-Score
PolitiFact	0.90	0.88	0.89	0.88
GossipCop	0.91	0.89	0.90	0.89
ReCOVery	0.92	0.90	0.91	0.90

4.3 Comparison with Baselines

4.3.1 Traditional Machine Learning Models

Baseline comparisons are conducted using traditional machine learning models, including Logistic Regression and Random Forest. These models rely on handcrafted features and do not capture complex relational dependencies within the network. As a result, their performance is comparatively lower when dealing with large-scale and dynamic social network data.

4.3.2 Static Network Models

Static network models analyze the structural properties of social networks without considering temporal evolution. While these approaches capture connectivity and community structure, they fail to model the dynamic nature of misinformation propagation. This limitation affects their ability to accurately predict future cascade behavior.

4.3.3 Proposed GNN-Based Model

The proposed graph-structured deep learning model outperforms all baseline approaches by effectively integrating structural and temporal information. The use of Graph Neural Networks enables the model to learn complex interaction patterns, while the temporal component captures evolving diffusion dynamics. As a result, the model achieves the highest accuracy, approximately 91%, along with improved precision, recall, and F1-score.

4.3.4 Comparative Performance Analysis

A detailed comparison of all models is presented in Table 3.

Table 3: Model Comparison

Model	Accuracy	Precision	Recall	F1-Score
Logistic Regression	0.81	0.79	0.76	0.77
Random Forest	0.85	0.83	0.82	0.82
Static Network Model	0.86	0.84	0.83	0.83
Proposed GNN Model	0.91	0.89	0.90	0.89

5. DISCUSSION

This section provides a comprehensive interpretation of the experimental findings and highlights the significance of the proposed temporal graph-based deep learning model in forecasting misinformation spread. It also explains the advantages of the model over existing approaches and discusses its practical implications in real-world scenarios.

5.1 Interpretation of Results

5.1.1 Analysis of Model Performance

The experimental results demonstrate that the proposed model achieves superior performance compared to traditional machine learning and static network-based approaches. The high accuracy, precision, recall, and F1-score indicate that the model is capable of effectively identifying misinformation patterns and predicting their propagation. This improvement can be attributed to the integration of structural and temporal features, which allows the model to capture both the relationships among users and the evolution of these relationships over time.

Furthermore, the model shows strong capability in forecasting cascade growth and predicting future interactions. By learning from historical propagation patterns, the system is able to estimate how misinformation spreads across the network and identify potential expansion points. The consistency of performance across multiple datasets suggests that the model is robust and generalizable to different types of misinformation scenarios.

5.2 Effectiveness of Temporal Graph Neural Networks

5.2.1 Why Temporal GNN Performs Better

The superior performance of the temporal Graph Neural Network can be explained by its ability to simultaneously model network structure and temporal dynamics. Traditional models either focus on static relationships or analyze data sequentially without considering relational dependencies. In contrast, the proposed approach combines graph-based learning with time-aware modeling, enabling it to capture complex interaction patterns that evolve over time.

The message passing mechanism in Graph Neural Networks allows nodes to aggregate information from their neighbors, thereby capturing local and global structural dependencies. When extended with temporal learning, the model can track how these dependencies change across different time intervals. This enables the detection of early-stage diffusion patterns that are often missed by static models. As a result, the temporal GNN provides a more accurate and comprehensive representation of misinformation propagation processes.

5.3 Real-World Implications

5.3.1 Early Warning Systems

One of the most significant implications of this research is its potential application in early warning systems for misinformation detection. The ability to forecast misinformation cascades at an early stage allows platforms to identify high-risk content before it spreads widely. By detecting early signals of rapid propagation, the model can support proactive intervention strategies, such as issuing warnings, limiting content visibility, or prioritizing fact-checking processes. This capability is crucial for reducing the societal impact of misinformation.

5.3.2 Platform Moderation and Policy Support

The proposed framework can also assist social media platforms in improving content moderation mechanisms. By analyzing interaction patterns and predicting future

propagation behavior, the model provides valuable insights into how misinformation spreads within communities. This information can be used to design targeted moderation strategies that focus on influential users or highly active communities. Additionally, policymakers can leverage these insights to develop evidence-based regulations and guidelines for managing misinformation in digital ecosystems. The integration of predictive analytics into moderation systems therefore represents a significant step toward creating more reliable and trustworthy online environments.

6. CONCLUSION

This study presents a comprehensive framework for forecasting misinformation spread in online social networks using graph-structured deep learning models. The research addresses the limitations of traditional approaches by integrating structural and temporal dynamics into a unified modeling framework. By representing social interactions as dynamic graphs and applying Graph Neural Networks with temporal learning mechanisms, the proposed model effectively captures complex propagation patterns of misinformation. The experimental results demonstrate that the model significantly outperforms traditional machine learning and static network-based methods in terms of accuracy, precision, recall, and F1-score. In particular, the ability to predict cascade growth and future interactions highlights the strength of the approach in understanding the evolution of misinformation over time. The incorporation of structural, temporal, and user-level features further enhances the model's predictive capability and robustness across different datasets. Moreover, the study emphasizes the importance of early detection of misinformation cascades, which is crucial for minimizing their societal impact. The findings confirm that combining graph-based learning with temporal modeling provides a more realistic and effective representation of information diffusion processes. Overall, this research contributes to the advancement of misinformation analysis by offering a scalable and efficient framework that can support real-time monitoring, early warning systems, and data-driven decision-making in online platforms.

7. FUTURE SCOPE OF RESEARCH

Future research can extend this work by incorporating multi-modal data sources, including textual, visual, and video content, to improve the robustness of misinformation detection. The integration of real-time streaming data can further enhance the model's capability for live monitoring and instant prediction. Additionally, exploring explainable artificial intelligence techniques can improve the interpretability of model predictions, making them more useful for decision-makers. Cross-platform analysis of misinformation spread and the development of scalable architectures for large-scale networks also represent promising directions for future investigation.

REFERENCES

1. Ahmed, H., Traore, I. and Saad, S. (2017) 'Detection of online fake news using n-gram analysis and machine learning techniques', International Conference on Intelligent, Secure, and Dependable Systems in Distributed and Cloud Environments, pp. 127–138.
2. Barabási, A.-L. (2016) Network Science. Cambridge: Cambridge University Press.
3. Castillo, C., Mendoza, M. and Poblete, B. (2011) 'Information credibility on Twitter', Proceedings of the 20th International Conference on World Wide Web, pp. 675–684.
4. Cheng, J., Adamic, L., Dow, P.A., Kleinberg, J. and Leskovec, J. (2014) 'Can cascades be predicted?', Proceedings of the 23rd International Conference on World Wide Web, pp. 925–936.
5. Cinelli, M., Quattrociocchi, W., Galeazzi, A., Valensise, C.M., Brugnoli, E., Schmidt, A.L., Zola, P., Zollo, F. and Scala, A. (2021) 'The echo chamber effect on social media', Proceedings of the National Academy of Sciences, 118(9), e2023301118.
6. Fawcett, T. (2006) 'An introduction to ROC analysis', Pattern Recognition Letters, 27(8), pp. 861–874.
7. Ferrara, E., Varol, O., Davis, C., Menczer, F. and Flammini, A. (2016) 'The rise of social bots', Communications of the ACM, 59(7), pp. 96–104.
8. Ferrara, E., Cresci, S. and Luceri, L. (2020) 'Misinformation, manipulation, and abuse on social media in the era of COVID-19', Journal of Computational Social Science, 3, pp. 271–277.
9. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O. and Dahl, G.E. (2017) 'Neural message passing for quantum chemistry', Proceedings of the 34th International Conference on Machine Learning, pp. 1263–1272.
10. Hamilton, W.L., Ying, R. and Leskovec, J. (2017) 'Inductive representation learning on large graphs', Advances in Neural Information Processing Systems, 30.
11. Holme, P. and Saramäki, J. (2012) 'Temporal networks', Physics Reports, 519(3), pp. 97–125.
12. Juul, J.L. and Ugander, J. (2021) 'Comparing information diffusion mechanisms by matching on cascade size', Proceedings of the National Academy of Sciences, 118(46).
13. Kazemi, S.M., Goel, R., Jain, K., Kobyzev, I., Sethi, A., Forsyth, P. and Poupart, P. (2020) 'Representation learning for dynamic graphs: A survey', Journal of Machine Learning Research, 21(70), pp. 1–73.
14. Kempe, D., Kleinberg, J. and Tardos, É. (2003) 'Maximizing the spread of influence through a social

network', Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 137–146.

15. Kumar, S. and Shah, N. (2018) 'False information on web and social media: A survey', arXiv preprint arXiv:1804.08559.
16. Lazer, D.M.J., Baum, M.A., Benkler, Y., Berinsky, A.J., Greenhill, K.M., Menczer, F., Metzger, M.J., Nyhan, B., Pennycook, G. and Rothschild, D. (2018) 'The science of fake news', Science, 359(6380), pp. 1094–1096.
17. LeCun, Y., Bengio, Y. and Hinton, G. (2015) 'Deep learning', Nature, 521(7553), pp. 436–444.
18. Lewandowsky, S., Ecker, U.K.H. and Cook, J. (2017) 'Beyond misinformation: Understanding and coping with the "post-truth" era', Journal of Applied Research in Memory and Cognition, 6(4), pp. 353–369.
19. Monti, F., Frasca, F., Eynard, D., Mannion, D. and Bronstein, M.M. (2019) 'Fake news detection on social media using geometric deep learning', arXiv preprint arXiv:1902.06673.
20. Newman, M. (2018) Networks. 2nd edn. Oxford: Oxford University Press.
21. Rossetti, G. and Cazabet, R. (2018) 'Community discovery in dynamic networks: A survey', ACM Computing Surveys, 51(2), pp. 1–37.
22. Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F. and Bronstein, M. (2020) 'Temporal graph networks for deep learning on dynamic graphs', arXiv preprint arXiv:2006.10637.
23. Shao, C., Ciampaglia, G.L., Varol, O., Flammini, A. and Menczer, F. (2018) 'The spread of fake news by social bots', Nature Communications, 9(1), 4787.
24. Shu, K., Sliva, A., Wang, S., Tang, J. and Liu, H. (2017) 'Fake news detection on social media: A data mining perspective', ACM SIGKDD Explorations Newsletter, 19(1), pp. 22–36.
25. Shu, K., Mahudeswaran, D. and Liu, H. (2020) 'FakeNewsNet: A data repository with news content, social context and dynamic information for studying fake news', Big Data, 8(3), pp. 171–188.
26. Sokolova, M. and Lapalme, G. (2009) 'A systematic analysis of performance measures for classification tasks', Information Processing & Management, 45(4), pp. 427–437.
27. Tandoc, E.C., Lim, Z.W. and Ling, R. (2018) 'Defining "fake news": A typology of scholarly definitions', Digital Journalism, 6(2), pp. 137–153.
28. Vosoughi, S., Roy, D. and Aral, S. (2018) 'The spread of true and false news online', Science, 359(6380), pp. 1146–1151.
29. Wardle, C. and Derakhshan, H. (2017) Information disorder: Toward an interdisciplinary framework for research and policy making. Council of Europe report.
30. Wu, L. and Liu, H. (2018) 'Tracing fake-news footprints: Characterizing social media messages by how they propagate', Proceedings of the 11th ACM International Conference on Web Search and Data Mining, pp. 637–645.
31. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C. and Philip, S.Y. (2021) 'A comprehensive survey on graph neural networks', IEEE Transactions on Neural Networks and Learning Systems, 32(1), pp. 4–24.
32. Zhou, X. and Zafarani, R. (2020) 'A survey of fake news: Fundamental theories, detection methods, and opportunities', ACM Computing Surveys, 53(5), pp. 1–40.
33. Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C. and Sun, M. (2020) 'Graph neural networks: A review of methods and applications', AI Open, 1, pp. 57–81.