

# DECOR: Degree-Corrected Social Graph Refinement for Fake News Detection

Jiaying Wu  
National University of Singapore  
jiayingwu@u.nus.edu

Bryan Hooi  
National University of Singapore  
bhooi@comp.nus.edu.sg

## ABSTRACT

Recent efforts in fake news detection have witnessed a surge of interest in using graph neural networks (GNNs) to exploit rich social context. Existing studies generally leverage fixed graph structures, assuming that the graphs accurately represent the related social engagements. However, *edge noise* remains a critical challenge in real-world graphs, as training on suboptimal structures can severely limit the expressiveness of GNNs. Despite initial efforts in graph structure learning (GSL), prior works often leverage node features to update edge weights, resulting in heavy computational costs that hinder the methods' applicability to large-scale social graphs. In this work, we approach the fake news detection problem with a novel aspect of *social graph refinement*. We find that the *degrees* of news article nodes exhibit distinctive patterns, which are indicative of news veracity. Guided by this, we propose DECOR, a novel application of Degree-Corrected Stochastic Blockmodels to the fake news detection problem. Specifically, we encapsulate our empirical observations into a lightweight social graph refinement component that iteratively updates the edge weights via a learnable degree correction mask, which allows for joint optimization with a GNN-based detector. Extensive experiments on two real-world benchmarks validate the effectiveness and efficiency of DECOR. <sup>1</sup>

## CCS CONCEPTS

• Information systems → Data mining; • Computing methodologies → Neural networks.

## KEYWORDS

Fake News; Graph Neural Networks; Social Network

### ACM Reference Format:

Jiaying Wu and Bryan Hooi. 2023. DECOR: Degree-Corrected Social Graph Refinement for Fake News Detection. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3580305.3599298>

## 1 INTRODUCTION

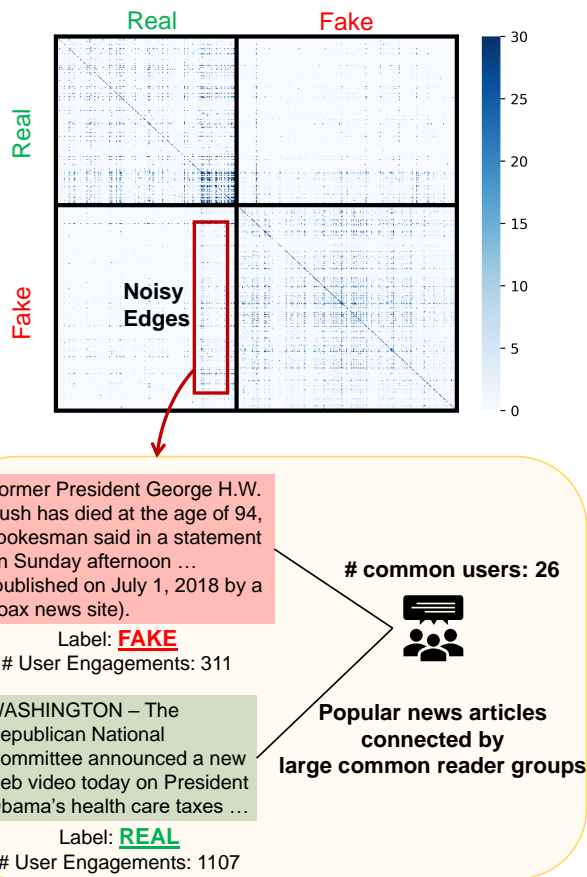
Automated detection of fake news stories containing intentionally distorted facts is a major focus of public discourse and scientific research [4, 5, 15, 30, 35]. Drawing inspiration from the expressive

<sup>1</sup>Data and code are available at: <https://github.com/jiayingwu19/DECOR>.



This work is licensed under a Creative Commons Attribution International 4.0 License.

KDD '23, August 6–10, 2023, Long Beach, CA, USA  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0103-0/23/08.  
<https://doi.org/10.1145/3580305.3599298>



**Figure 1: A motivating example for social graph refinement. Darker spots represent a larger number of common readers between two news articles. Weights are clipped at 30 for a clear visualization.**

Graph Neural Networks (GNNs) [14, 19, 46], a substantial body of research incorporates graphs with rich social context [22, 26, 37] to encode the news dissemination patterns and user responses. Despite varying choices of feature types and GNN backbones, existing approaches are relatively consistent in the design of social graphs. More specifically, the graphs typically contain social users and news articles, which allow GNNs to leverage the relations between structural patterns and news veracity (e.g., closely connected nodes tend to have similar preferences or veracity-related properties). This facilitates the aggregation and propagation of crowd wisdom among connected articles and users, yielding more accurate predictions. Therefore, existing works typically consider social graphs as a high

fidelity representation of the social context, the structure of which is kept unchanged throughout model training.

However, we find that *noisy edges* remain an inevitable challenge for learning on social graphs. A prominent cause is that influential news articles tend to share a large number of common readers (i.e. social users), regardless of their veracity labels. In Figure 1, we illustrate one such case via visualizing the adjacency matrix of a graph constructed with news articles, termed as the news engagement graph (detailed formulation relegated to Section 3.2). Here, the edge weights are positively correlated with the number of common readers, and larger edge weights imply closer connections between news articles. Figure 1 shows that the largest weights are assigned to the diagonal, and areas representing real news pairs and fake news pairs are also darker. This is expected, given the self-loops and frequent interactions within groups of real news and fake news spreaders. However, the figure also shows scattered dark spots representing noisy edges between real news and fake news. We observe that **edge noise is degree-related**, in that large edge weights are often distributed along certain rows and columns.

Noisy edges severely undermine the effectiveness of GNNs, as the message passing mechanism [13] propagates noise and contaminates node representations. However, little effort has been made to mitigate this issue in the fake news detection scenario. Despite preliminary efforts of graph structure learning (GSL) methods in denoising edges for real-world graphs (e.g. citation networks) [6, 16], existing GSL methods cannot be readily applied to fake news detection, as they generally leverage pairwise node feature similarity to guide edge weight updates. Given the large scale of social graphs, similarity-guided GSL becomes less feasible and raises critical deployment challenges.

In this work, we investigate the fake news detection problem from a novel aspect of **social graph refinement**. Given a set of news articles, we construct and refine a news engagement graph that connects the articles with common readers. Guided by our observation of degree-related edge noise, we explore veracity-related degree patterns on the news readership graph, and make two key findings: (1) nodes representing fake news and real news exhibit distinctive degree distributions; and (2) grouping edges by the veracity labels of the articles they connect, different edge groups demonstrate a clear difference regarding the relationship between degrees and the number of common readers.

Motivated by our empirical findings on veracity-related degree and co-engagement patterns, we present **Degree-Corrected Social Graph Refinement (DECOR)**, a novel social graph refinement framework for fake news detection. DECOR is based on a flexible extension of the Degree-Corrected Stochastic Blockmodel (DCSBM) [17], a graph generative model that allows us to simultaneously consider the effects of degree and node labels, in a tractable probabilistic manner. DECOR suppresses noise in the news engagement graph by downweighting the noisy edges, specifically via learning a *social degree correction mask* based on a theoretically motivated likelihood ratio-based statistic under the DCSBM model, with a nonlinear relaxation to improve the flexibility of the model. DECOR utilizes the degree correction mask to adjust the edge weights of the news engagement graph, which is then jointly optimized with a GNN-based classifier to predict news veracity.

In summary, our contributions are as follows:

- **Empirical Findings:** We present two novel findings, on how both *degree* and *co-engagement* closely relate to news veracity.
- **Principled DCSBM-based GSL:** Motivated by our empirical findings, we propose DECOR, a GSL approach for reducing edge noise, based on a theoretically motivated *likelihood ratio*-based statistic under the DCSBM model, combined with a nonlinear relaxation.
- **Efficiency:** Unlike existing GSL approaches, DECOR avoids using high dimensional features as input for GSL, and is also linear in the number of edges, thus being 7.6 - 34.1 times faster than existing GSL approaches.
- **Effectiveness:** DECOR improves F1 score by 4.55% and 2.51% compared to the best baseline on two real-world fake news detection benchmarks, consistently improves the performance of multiple GNN baselines in a plug-and-play manner, and outperforms baselines under label scarcity.

## 2 RELATED WORK

### 2.1 Fake News Detection

Fake news detection is commonly considered as a binary classification problem, with the goal of accurately predicting a given news article as real or fake. Among existing studies, **content-based methods** extract semantic patterns from the news content using a wide range of deep learning architectures that include RNNs [29] and pre-trained language models (PLMs) [20, 28]. Some methods also guide model prediction with auxiliary information including knowledge bases [5, 10, 15, 45], evidence from external sources [3, 34, 47], visual information [2, 32, 42, 51], and signals from the news environment [33]. As fake news detection is often deeply rooted in the social context, **propagation-based methods** incorporate various social features including user responses and opinions [24, 30, 35, 37, 48, 49], user-user following relations [22], news sources [26], and user history posts [9] to guide model prediction. Despite the rich social information incorporated, little effort has been made to explore direct relations between news articles and the properties of veracity-related news-news connections. Moreover, many methods are vulnerable to structural noise in social graphs, as they typically adopt fixed graph structures during training.

### 2.2 Structure Learning for Robust GNNs

Graph Neural Networks (GNNs) have demonstrated impressive potential in learning node and graph representations [14, 19, 25, 46]. Despite the prior success, extensive studies have demonstrated that GNNs are highly vulnerable to adversarial attacks in terms of structural noise [7, 41, 52]. To alleviate this issue, numerous works have focused on learning optimized structures for real-world graphs, specifically via edge denoising [6, 11, 16, 38, 44]. Motivated by the observation that noisy edges connect nodes with dissimilar features [11], existing methods are generally guided by feature similarity measures. For instance, [44] conducts edge pruning based on the Jaccard similarity between paired node features, Pro-GNN [16] employs the feature smoothness regularization alongside low-rank constraints, and RS-GNN [6] utilizes node feature similarity to guide the link prediction process. Nevertheless, graph structure learning (GSL) remains underexplored under the social context of

fake news detection. Existing GSL methods are not readily applicable to this task, given the high computational costs incurred in computing pairwise similarity measures between high-dimensional news article representations on large-scale social graphs. While initial efforts have been made in conditioning the edge metric with node degrees for coordination detection [50], the fixed adjustment formula adopted by existing work cannot fully capture the complex relations between degree-related properties, which may vary greatly across datasets. To the best of our knowledge, we propose the first learnable framework for social graph refinement, which leverages low-dimensional degree-related properties to flexibly adjust the edge weights of a news engagement graph for enhanced fake news detection.

### 3 PRELIMINARY ANALYSIS

In this section, we formally define the fake news detection problem, establish a social context graph that encodes user engagements in disseminating news articles, and conduct preliminary analysis to explore the veracity-related structural patterns.

#### 3.1 Problem Formulation

Let  $\mathcal{D}$  be a fake news detection dataset containing  $N$  samples. In the social media setting, we define the dataset as

$$\mathcal{D} = \{\mathcal{P}, \mathcal{U}, \mathcal{R}\},$$


where  $\mathcal{P} = \{p_1, p_2, \dots, p_N\}$  is a set of questionable **news articles**,  $\mathcal{U} = \{u_1, u_2, \dots\}$  is a set of related **social users** who have spread at least one article in  $\mathcal{P}$  via reposting on social media.  $\mathcal{R}$  represents the set of **social user engagements**, in which  $r \in \mathcal{R}$  is defined as a triple  $\{(u, p, k) | u \in \mathcal{U}, p \in \mathcal{P}\}$  (i.e. user  $u$  has given  $k$  responses to the news article  $p$  in terms of *reposts*). In line with most existing studies, we treat fake news detection on social media as a binary classification problem. Specifically,  $\mathcal{P}$  is split into training set  $\mathcal{P}_{train}$  and test set  $\mathcal{P}_{test}$ . Article  $p \in \mathcal{P}_{train}$  is associated with a ground-truth label  $y$  of 1 if  $p$  is fake, and 0 otherwise. We formulate the problem as follows:

**PROBLEM 1 (FAKE NEWS DETECTION ON SOCIAL MEDIA).** *Given a news dataset  $\mathcal{D} = \{\mathcal{P}, \mathcal{U}, \mathcal{R}\}$  and ground-truth training labels  $\mathcal{Y}_{train}$ , the goal is to learn a classifier  $f$  that, given test articles  $\mathcal{P}_{test}$ , is able to predict the corresponding veracity labels  $\mathcal{Y}_{test}$ .*

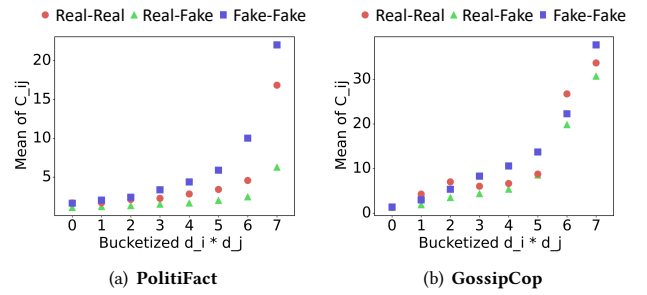
#### 3.2 News Engagement Graph

The positive correlation between social user preferences and the user’s news consumption habits has been acknowledged by prior research [1]. Specifically, social media creates an *echo chamber*, where individual beliefs can be continuously reinforced by communication and repetition within like-minded social groups [12].

Motivated by this, we propose to capture the news veracity signals embedded in social user engagements. To distill a comprehensive representation of user preferences, we set a threshold to filter the users with less than 3 engagements with news articles, and focus on a subset  $\mathcal{U}_A \subset \mathcal{U}$  containing active users. Specifically, we construct a *user engagement matrix*  $\mathbf{E} \in \mathbb{R}^{|\mathcal{U}_A| \times N}$ . Element  $\mathbf{E}_{ij}$  represents the number of interactions between user  $u_i$  and news article  $p_j$ , the value of which is retrieved from the corresponding entry  $(u_i, p_j, k_{ij}) \in \mathcal{R}$ .



**Figure 2: KDE plot of node degree distributions on the news engagement graph.**



**Figure 3: News co-engagement patterns of news article pairs. Edges in  $\mathcal{G}$  represent shared readership between articles, and are grouped based on the articles’ veracity labels.**

Given the news consumption patterns of active social users, we further propose to link the news articles that attract similar user groups via constructing an weighted undirected *news engagement graph*  $\mathcal{G} = \{\mathcal{P}, \mathcal{E}\}$ . The adjacency matrix  $\mathbf{A} \in \mathbb{R}^{N \times N}$  of  $\mathcal{G}$  is formulated based on overlapping user engagement patterns in  $\mathbf{E}$ , specifically as:

$$\mathbf{A} = \mathbf{E}^T \mathbf{E}. \quad (1)$$

Intuitively, element  $\mathbf{A}_{nk}$  in  $\mathbf{A}$  can be interpreted as the number of 2-hop paths (i.e., news - user - news) between two news articles  $p_n$  and  $p_k$ . Hence, a larger  $\mathbf{A}_{nk}$  value represents stronger common interest between the reader groups of news article, implying shared opinions or beliefs in the users’ news consumption preferences.

#### 3.3 Empirical Observations

In this subsection, we conduct preliminary analysis on real-world news to explore the veracity-related structural properties on the news engagement graph. We observe that fake and real news exhibit distinctive patterns in terms of weighted node **degrees**, motivated by which we design a degree-based social graph refinement framework to mitigate the edge noise issue in Section 4. Our analysis is based on the FakeNewsNet [36] benchmark, which consists of the PolitiFact and GossipCop datasets.

**3.3.1 Degree-Veracity Correlations.** We first explore how the degree of a news article node is related to its veracity label. In other

words, **do fake news articles attract more or less user engagements than real news?**

Recall that we have a news engagement graph  $\mathcal{G} = \{\mathcal{P}, \mathcal{E}\}$  with adjacency matrix  $A$ . The weighted node degrees in  $A$  can be used to measure the intensity of user engagements for each news article. In Figure 2, we visualize the degree distributions of fake and real news with a kernel distribution estimation (KDE) plot, which depicts the node degrees with a continuous probability density curve. We make the following observation:

**OBSERVATION 1.** *On the news engagement graph, the degree distributions of nodes representing fake and real news articles show a clear difference. Note that different datasets can exhibit varying domain-specific patterns; for instance, in the GossipCop dataset containing celebrity news, real news tend to attract more engagements from active social users. However, this pattern does not apply to the politics-related PolitiFact dataset.*

**3.3.2 News Co-Engagement.** Next, we explore the degree-related properties of news article pairs connected by common readers (i.e. active social users in  $\mathcal{U}_A$ ). Intuitively, given a pair of news articles  $p_i$  and  $p_j$  that share at least 1 reader, the corresponding edge  $e_{ij} \in \mathcal{E}$  in the news engagement graph can be divided into three groups according to the veracity labels of  $p_i$  and  $p_j$ : (1) real news pairs; (2) real-fake pairs; and (3) fake news pairs.

To quantify the shared user engagements between news article nodes w.r.t. the corresponding degrees, we compute a “co-engagement” score  $C_{ij}$  for news articles  $p_i$  and  $p_j$ , formulated as:

**DEFINITION 1 (NEWS CO-ENGAGEMENT).**

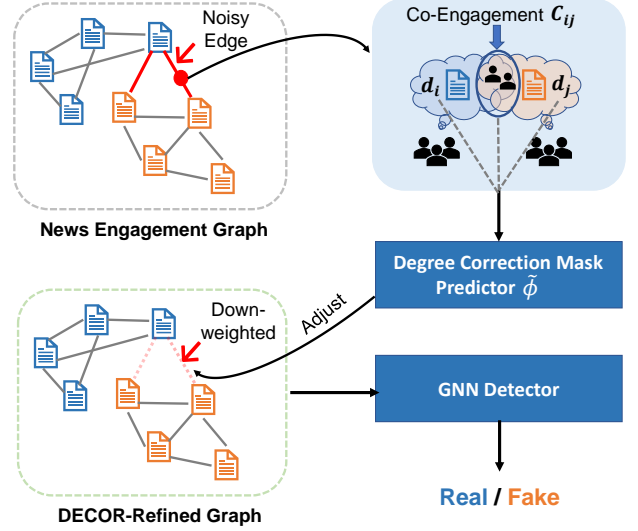
$$C_{ij} = |\mathcal{U}_i \cap \mathcal{U}_j|,$$

where  $\mathcal{U}_i \subset \mathcal{U}$  and  $\mathcal{U}_j \subset \mathcal{U}$  are the sets of social users that engage with  $p_i$  and  $p_j$ , respectively.

We investigate the following question: **given an edge, are there any associations between its group, and the news co-engagement of the two nodes it connects?** In Figure 3, we bucketize the edges by the value of  $d_i \times d_j$ , and plot the news co-engagement scores w.r.t. the edge groups. Note that here we adopt the product of degrees to distinguish edges with high values for both  $d_i$  and  $d_j$ , and also motivated by our theoretical results in Section 4.1. Across the buckets, we observe the following pattern on news co-engagement:

**OBSERVATION 2.** *Given the degrees, fake news pairs tend to have higher  $C_{ij}$  (i.e. more common users than expected given the degrees), while real-fake pairs tend to have lower  $C_{ij}$  than both real news pairs and fake news pairs.*

Our two empirical observations provide distinctive degree-related cues pertaining to nodes (i.e. news articles) and edges (i.e. user engagements) on the news engagement graph (extended analysis and discussion are relegated to Appendix A). These patterns can guide a model in suppressing the noisy edges, as they can be leveraged to identify which edges are more likely to connect news articles of the same veracity. Meanwhile, we find that differences in the degree distributions can be complex (e.g., as shown in Figure 2, fake news attract more user engagements than real news in PolitiFact, but less in GossipCop). This motivates our following degree-based innovations for a learnable social graph refinement approach.



**Figure 4: Overview of the proposed Degree-Corrected Social Graph Refinement (DECOR) framework.**

## 4 PROPOSED FRAMEWORK – DECOR

Motivated by our empirical findings on veracity-related degree patterns, we propose the DECOR framework for degree-corrected social graph refinement (overviewed in Figure 4). DECOR can be considered as a novel extension of the Degree-Corrected Stochastic Blockmodel (DCSBM) [17] to the fake news detection scenario, which empowers fake news detectors with effective denoising of user engagements. Given a pair of news articles connected by common users, we propose a social degree correction module to adjust the corresponding edge weight using degrees and the news co-engagement. This module is jointly optimized with the GNN classifier, which leverages the corrected edge weights and news article features to predict the news veracity labels.

### 4.1 Connection with the DCSBM Model

In Section 3.3, we observed that degree patterns are closely related to news veracity labels. Next, we formally demonstrate these connections from a theoretical perspective based on the DCSBM model [17], a generative model for graphs that derives edge placement likelihoods in a degree-based manner. The benefit of DCSBM is that it allows us to simultaneously model the effect of *degree patterns* and *class labels*, which are of key interest, in a tractable probabilistic way. Based on the DCSBM model, we will then theoretically derive a principled likelihood ratio-based approach for graph structure learning for the fake news detection application.

**Framework.** We first formulate the standard DCSBM under our fake news detection scenario. Recall the news engagement graph  $\mathcal{G} = \{\mathcal{P}, \mathcal{E}\}$  formulated in Section 3.2, where  $|\mathcal{P}| = N$ . Each news article node in  $\mathcal{G}$  is associated with a class label from the label space  $\mathcal{Z} = \{0, 1\}$ . Consider a pair of news article nodes  $p_i \in \mathcal{P}$  and  $p_j \in \mathcal{P}$  with co-engagement  $C_{ij}$ . The nodes have class labels  $z_i \in \mathcal{Z}$  and  $z_j \in \mathcal{Z}$ , respectively. Recall that  $C_{ij}$  is defined as the number of common users between  $p_i$  and  $p_j$ .

Next, to formulate structure learning under the DCSBM model, our basic intuition is that *same-class edges* (i.e., edges  $e_{ij}$  where  $z_i = z_j$ ) are more likely to be useful and informative than *cross-class edges* (i.e., edges where  $z_i \neq z_j$ ), and hence, structure learning should aim to give a higher weight to same-class edges. Intuitively, cross-class edges tend to indicate noisy edges, as in the example in Figure 1, where the co-engagement between them arises just by chance. Moreover, since our main goal is to classify  $p_i$ , identifying edges where  $z_i = z_j$  clearly provides highly useful information for this task. Hence, our key idea is to perform structure learning by deriving the *same-class likelihood ratio*:

**DEFINITION 2 (SAME-CLASS LIKELIHOOD RATIO).** *The same-class likelihood ratio, i.e. the likelihood ratio for  $z_i = z_j$  over  $z_i \neq z_j$  when observing  $C_{ij}$  edges between  $p_i$  and  $p_j$ , is*

$$LR_{ij} := \frac{\mathbb{P}(C_{ij}|z_i = z_j)}{\mathbb{P}(C_{ij}|z_i \neq z_j)}. \quad (2)$$

The higher this likelihood ratio, the more evidence the data (specifically,  $C_{ij}$ ) gives in favor of  $z_i = z_j$  over  $z_i \neq z_j$ ; and hence, structure learning should give a higher weight to such edges.

*Derivation.* Under the DCSBM model, the  $C_{ij}$  edges between  $p_i$  and  $p_j$  are independently Poisson distributed, i.e.,  $C_{ij} \sim \text{Poi}(\lambda_{ij})$ , where  $\lambda_{ij}$  denotes the expected number of edges:

$$\lambda_{ij} = \begin{cases} \beta_i \beta_j p & \text{if } z_i = z_j \\ \beta_i \beta_j q & \text{if } z_i \neq z_j \end{cases}, \quad (3)$$

where  $\beta_i$  and  $\beta_j$  are the ‘‘degree correction parameters’’ that allow us to generate nodes with different degrees.  $p$  and  $q$  are parameters controlling the rate at which edges are generated under the same-class and cross-class cases, respectively. Generally, we have  $p > q$ , i.e., same-class edges have a higher tendency to be generated.

The corresponding maximum likelihood values  $\hat{\beta}_i$  and  $\hat{\beta}_j$  for  $\beta_i$  and  $\beta_j$  are given as

$$\hat{\beta}_i = \frac{d_i}{m}, \quad \hat{\beta}_j = \frac{d_j}{m}, \quad (4)$$

in the DCSBM model [17], where  $m = |\mathcal{E}|$  denotes the number of edges.  $d_i$  and  $d_j$  respectively refer to the weighted degrees of nodes  $p_i$  and  $p_j$ .

Since  $C_{ij} \sim \text{Poi}(\lambda_{ij})$ , the likelihood ratio  $LR_{ij}$  for  $z_i = z_j$  over  $z_i \neq z_j$  can be derived as:

$$\begin{aligned} LR_{ij} &= \frac{\mathbb{P}(C_{ij}|z_i = z_j)}{\mathbb{P}(C_{ij}|z_i \neq z_j)} \\ &= \frac{e^{-\beta_i \beta_j p} (\beta_i \beta_j p)^{C_{ij}}}{e^{-\beta_i \beta_j q} (\beta_i \beta_j q)^{C_{ij}}} \\ &= e^{-\beta_i \beta_j (p-q)} \left(\frac{p}{q}\right)^{C_{ij}}. \end{aligned} \quad (5)$$

Substituting the  $\hat{\beta}_i$  and  $\hat{\beta}_j$  given in Eq.4 into Eq.5, we derive the maximum likelihood estimate for  $LR_{ij}$ :

$$\text{MLE}(LR_{ij}) = e^{-\frac{d_i d_j}{m^2} (p-q)} \left(\frac{p}{q}\right)^{C_{ij}}. \quad (6)$$

Treating  $m, p, q$  as fixed (since they are shared by all nodes), we thus see that the MLE is a function of  $C_{ij}, d_i$  and  $d_j$ : in particular,

it is a *log-linear* function of  $C_{ij}$  and  $d_i d_j$ :

$$\text{MLE}(LR_{ij}) = \Phi(C_{ij}, d_i, d_j) := e^{-\frac{d_i d_j}{m^2} (p-q)} \left(\frac{p}{q}\right)^{C_{ij}} \quad (7)$$

$$= \exp \left[ \left( \frac{C_{ij}}{d_i d_j} \right) \cdot \left( \log(p) - \log(q) - \frac{p-q}{m^2} \right) \right] \quad (8)$$

*Implications.* We first note that Eq. 8 agrees with our empirical finding in Observation 2: if we fix  $d_i d_j$  in Eq. 8, then as long as  $\log(p) - \log(q) > 0$ , we observe that higher  $C_{ij}$  is associated with a higher  $LR_{ij}$ , and thus a higher probability of same-class edges ( $z_i = z_j$ ), agreeing with Figure 3 where the Real-Fake edges have lowest  $C_{ij}$  for a given  $d_i d_j$ .

For structure learning purposes, we could simply use  $\Phi(C_{ij}, d_i, d_j)$ , which we recall is an estimator for  $LR_{ij} = \frac{\mathbb{P}(C_{ij}|z_i=z_j)}{\mathbb{P}(C_{ij}|z_i \neq z_j)}$ . However, the standard DCSBM model is built upon relatively strong assumptions (e.g. pre-defined  $p$  and  $q$  values); for fitting real data, we would like to relax these assumptions and allow the model to be flexibly learned from data. The DCSBM model contains very few learnable parameters, which is a fundamental limitation in adapting to the complex degree-based patterns in the news engagement graph. This motivates us to develop DECOR, a degree-based learnable social graph refinement framework, which we will next describe in detail, by *relaxing the assumption of log-linearity*: that is, instead of treating  $\Phi(C_{ij}, d_i, d_j)$  as a fixed and log-linear function defined in Eq. 8, we instead treat it as a *learnable non-linear function*  $\tilde{\Phi}(C_{ij}, d_i, d_j)$  to be updated jointly with the rest of the model, during the structure learning process.

## 4.2 Social Degree Correction

As illustrated in Figure 1, the news engagement graph contains structural noise. In light of our empirical findings on degree-verity relationships and the DCSBM framework, we propose to learn a degree-corrected social graph that downweights the noisy edges to eliminate their negative impacts and facilitate fake news detection via GNN-based classifiers.

Recall that the type of an edge in the news engagement graph (i.e. connecting new articles of same or different veracity) is characterized by the co-engagement and degrees of the connected articles. Motivated by the DCSBM model’s degree-based probabilistic derivation of edge placement likelihood, we propose to adjust edge weights in the news engagement graph via learning a *social degree correction mask*  $\mathbf{M} \in \mathbb{R}^{N \times N}$ , where  $M_{ij}$  in the interval  $(0, 1)$  represents the degree correction score for edge  $e_{ij}$  between news article nodes  $p_i$  and  $p_j$ .

The value of  $M_{ij}$  is predicted given co-engagement  $C_{ij}$  of articles  $p_i$  and  $p_j$ , and the articles’ weighted node degrees  $d_i$  and  $d_j$  from the news engagement graph. Specifically, we adopt a neural predictor to obtain  $s_{ij} \in \mathbb{R}^2$ , which contains two scores for edge preservation and elimination, respectively:

$$s_{ij} = \tilde{\Phi}(C_{ij}, d_i, d_j). \quad (9)$$

$\tilde{\Phi}(\cdot)$  is a MLP-based architecture, and can be considered as a learnable extension of Eq.8 in the DCSBM model.

The scores in  $s_{ij}$  are normalized via the softmax function. To preserve computational efficiency, we design the social degree correction process as *pruning*. In other words, we conduct Eq.9 on all

the news pairs connected by common users to obtain the corresponding degree correction scores:

$$\mathbf{M}_{ij} = \begin{cases} v_{ij} & \text{if } C_{ij} \neq 0 \\ 0 & \text{else} \end{cases}. \quad (10)$$

where  $v_{ij}$  denotes the softmax-normalized score in  $s_{ij}$  that correlates with edge preservation.

Given the co-engagement matrix  $\mathbf{C}$  of news engagement graph  $\mathcal{G}$ , we utilize  $\mathbf{M}$  to obtain a degree-corrected adjacency matrix  $\mathbf{A}_c$ :

$$\hat{\mathbf{A}} = \mathbf{C} \cdot \mathbf{M} + \mathbf{I} \quad (11)$$

$$\mathbf{A}_c = \mathbf{D}^{-\frac{1}{2}} \hat{\mathbf{A}} \mathbf{D}^{-\frac{1}{2}}, \quad (12)$$

where  $\mathbf{I}$  represents an identity matrix of size  $N$ , and  $\mathbf{D}$  is the diagonal matrix of degrees for  $\hat{\mathbf{A}}$ .

Through the above operations, noisy edges in the news engagement graph are assigned smaller weights, as  $\tilde{\Phi}(\cdot)$  in Eq.9 leverages degree-based properties to predict a low degree correction score.

### 4.3 Prediction on Degree-Corrected Graph

With the degree-corrected adjacency matrix  $\mathbf{A}_c$ , we can leverage the powerful GNN architectures (e.g. GCN [19], GIN [46] and GraphConv [25]) to predict the veracity labels of article nodes in the degree-corrected news engagement graph.

Central to GNNs is the message-passing mechanism [13], which follows an iterative scheme of updating node representations based on information aggregation among the node’s neighborhood. For a news article  $p \in \mathcal{P}$ , the initial news article feature  $\mathbf{h}_p^{(0)}$  is set as the news content representation  $\mathbf{x}_p$ :

$$\mathbf{h}_p^{(0)} = \mathbf{x}_p, \quad (13)$$

where  $\mathbf{x}_p$  is extracted from news article  $p$  via a pre-trained language model  $\mathcal{M}$  with frozen parameters. At the  $k$ -th layer of a GNN, the news article representation  $\mathbf{h}_p^{(k)}$  is obtained via:

$$\mathbf{m}_p^{(k)} = \text{AGGREGATE}^{(k)} \left( \left\{ \mathbf{h}_u^{(k-1)}, \forall u \in \mathcal{N}(p) \right\} \right) \quad (14)$$

$$\mathbf{h}_p^{(k)} = \text{COMBINE}^{(k)} \left( \mathbf{h}_p^{(k-1)}, \mathbf{m}_p^{(k)} \right), \quad (15)$$

where  $\mathcal{N}(p)$  denotes the neighbors of  $p$  on the news engagement graph, and  $\mathbf{m}_p^{(k)}$  is the aggregated information from  $\mathcal{N}(p)$ .

Let  $\mathbf{h}_p \in \mathbb{R}^2$  be the output of the GNN-based classifier for node  $p$ . Then, the news veracity label of  $p$  is predicted as  $\tilde{y}_p = \text{softmax}(\mathbf{h}_p)$ . During training, we minimize the following cross entropy loss:

$$\mathcal{L} = \sum_{p \in \mathcal{P}_{train}} \text{CELoss}(\tilde{y}_p, y_p). \quad (16)$$

The degree correction mask predictor  $\tilde{\Phi}(\cdot)$  is jointly optimized with the GNN-based classifier. DECOR utilizes low-dimensional degree-related properties to guide the social degree correction operations, which facilitates edge denoising on  $\mathcal{G}$  without loss of computational efficiency.

**Table 1: Dataset statistics.**

Dataset	PolitiFact	GossipCop
# News Articles	497	16,599
# Real News	225	12,641
# Fake News	272	3,958
# User-News Engagements	227,184	963,009
# Distinct Users	143,481	202,907

## 5 EXPERIMENTS

In this section, we empirically evaluate DECOR to answer the following five research questions:

- **Fake News Detection Performance** (Section 5.2): How well does DECOR perform compared with competitive baselines?
- **Ablation Study** (Section 5.3): How effective are co-engagement and degree patterns, respectively, in improving the fake news detection performance of DECOR?
- **Limited Training Data** (Section 5.4): Does DECOR perform well under label sparsity?
- **Computational Efficiency** (Section 5.5): How efficient is DECOR compared with existing GSL methods?
- **Case Study** (Section 5.6): Does DECOR downweight the noisy edges connecting influential real and fake news articles?

### 5.1 Experimental Setup

**5.1.1 Datasets.** We evaluate DECOR on the public benchmark FakeNewsNet [36], which consists of two real-world datasets: PolitiFact and GossipCop. Both datasets contain news articles annotated by leading fact-checking websites and the articles’ related social user engagements from Twitter. The descriptive statistics of the datasets are summarized in Table 1.

To simulate the real-world scenarios, we split the news samples following a *temporal* order. Specifically, the most recent 20% real and fake news instances constitute the test set, and the remaining 80% instances posted earlier serve as the training set.

**5.1.2 Baselines.** We benchmark DECOR against twelve representative baseline methods, which can be categorized into the following three groups by model architecture:

**News content based methods (G1)** leverage the semantic features in the news articles. Specifically, **dFEND** [35] is a content-based variant of dFEND [35] without incorporating user comment texts, which utilizes a hierarchical network with the co-attention mechanism. **SAFE** is a content-based variant of SAFE [51] without incorporating visual information from images, which leverages a CNN-based fake news detector. **SentGCN** [39] models each news article as a graph of sentences, and utilize the GCN [19] architecture for news veracity prediction. **BERT** [8] and **DistilBERT** [31] (with model names BERT-base and DistilBERT-base, respectively) are large pre-trained bidirectional Transformers, which we fine-tune to the downstream task of fake news detection.

**Social graph based methods (G2)** encode the social context into graph structures, and leverage GNNs to learn news article representations. Specifically, **GCNFN** [24] leverages user responses and user following relations to construct a propagation tree for

**Table 2: Performance comparison between DECOR and baselines (G1: Content-based, G2: Graph-based, and G3: GSL-based). Bold and underline indicates the best overall and baseline performance, respectively. \* denotes that DECOR performs significantly better than the corresponding GNN backbone at  $p < 0.01$  level using the Wilcoxon signed-rank test.**

Method	PolitiFact				GossipCop				
	Acc.	Prec.	Rec.	F1	Acc.	Prec.	Rec.	F1	
<b>G1</b>	dFEND\c [35]	0.8207	0.8226	0.8198	0.8195	0.7998	0.7251	0.6787	0.6936
	SAFEv [51]	0.8146	0.8185	0.8114	0.8121	0.7920	0.7090	0.6739	0.6866
	SentGCN [39]	0.8404	0.8567	0.8446	0.8385	0.7935	0.7130	0.6797	0.6913
	BERT [8]	0.8586	0.8692	0.8584	0.8564	0.8498	0.8084	0.7480	0.7703
	DistilBERT [31]	0.8419	0.8518	0.8426	0.8403	0.8349	0.7765	0.7422	0.7563
<b>G2</b>	GCNFN [24]	0.8687	0.8694	0.8680	0.8674	0.8520	0.9016	0.6950	0.7354
	FANG [26]	0.8601	0.8625	0.8647	0.8599	0.8715	0.8835	0.7506	0.7897
	GCN [19]	0.8965	0.9027	0.9034	0.8964	0.9146	<u>0.9370</u>	0.8281	0.8668
	GIN [46]	<u>0.9025</u>	<u>0.9037</u>	<u>0.9068</u>	<u>0.9024</u>	<u>0.9197</u>	0.9308	<u>0.8450</u>	<u>0.8780</u>
	GraphConv [25]	0.8990	0.9007	0.9037	0.8989	0.9150	0.9172	0.8426	0.8720
<b>G3</b>	Pro-GNN [16]	0.7747	0.7779	0.7690	0.7691	OOM			
	RS-GNN [6]	0.7439	0.7839	0.7253	0.7210	OOM			
<b>Ours</b>	GCN w/ DECOR	<b>0.9480*</b>	<b>0.9483*</b>	<b>0.9520*</b>	<b>0.9479*</b>	0.9310*	<b>0.9377</b>	0.8694*	0.8973*
	GIN w/ DECOR	0.9399*	0.9398*	0.9433*	0.9398*	<b>0.9333*</b>	0.9294	<b>0.8828*</b>	<b>0.9031*</b>
	GraphConv w/ DECOR	0.9379*	0.9376*	0.9412*	0.9377*	0.9259*	0.9130	0.8805*	0.8945*

each news article. **FANG** [26] establishes a comprehensive social graph with users, news and sources, and learns the representations with GraphSAGE [14]. We also apply three representative GNN architectures on our proposed news engagement graph, namely **GCN** [19], **GIN** [46], and **GraphConv** [25]. For a fair comparison, we only implement the model components involving news articles, social user identities, and user-news relations.

**Graph Structure learning (GSL) methods (G3)** aim to enhance representation via learning an optimized graph structure. We implement two GSL methods that focus on edge denoising, **Pro-GNN** [16] applies low-rank and sparsity properties to learn a clean graph structure that is similar to the original graph. **RS-GNN** [6] simultaneously learns a denoised graph and a robust GNN via constructing a link predictor guided by node feature similarity.

**5.1.3 Evaluation Metrics.** Following prior works [35, 51], we adopt four widely-used metrics to evaluate the performance of fake news detection methods: Accuracy (**Acc.**), Precision (**Prec.**), Recall (**Rec.**) and F1 Score (**F1**). In all experiments, we report the average metrics across 20 different runs of each method.

**5.1.4 Implementation Details.** We implement our proposed DECOR model and its variants based on PyTorch 1.10.0 with CUDA 11.1, and train them on a server running Ubuntu 18.04 with NVIDIA RTX 3090 GPU and Intel(R) Xeon(R) Gold 6226R CPU @ 2.90GHz. To construct DECOR’s news engagement graph, we select active social users with least 3 reposts, and threshold a user’s maximum number of interactions with each news article at 1% of the total number of news articles. We extract 768-dimensional news article features via a pre-trained BERT model with frozen parameters; specifically, we utilize pre-trained weights from HuggingFace Transformers 4.13.0 [43] (model name: bert-base-uncased). The predictor  $\tilde{\Phi}(\cdot)$  for social degree correction is a 2-layer MLP with hidden size 16 for PolitiFact

and 8 for GossipCop. The GNN architecture is set to 2 layers with 64 hidden dimensions. The model is trained for 800 epochs, and model parameters are updated for via an Adam optimizer [18] with learning rate 0.0005.

Technically, our framework is model-agnostic, which could coordinate with various GNN models on the news engagement graph. Here, we select three representative GNN architectures as backbones: GCN [19], GIN [46] and GraphConv [25]. For the implementation of baseline methods, we follow the architectures and hyperparameter values suggested by their respective authors.

## 5.2 Performance Comparison

This subsection compares DECOR with various content-based, graph-based and GSL baselines on fake news detection.

Table 2 shows that DECOR consistently outperforms the competitive baseline methods by significant margins ( $p < 0.01$ ). We make the following observations: **(1)** Among the five content-based methods (G1), pre-trained language models (PLMs) outperform the “train-from-scratch” methods. The effectiveness of PLMs demonstrates the benefits of pre-training on large-scale corpora, from which the model obtains rich semantic knowledge. **(2)** Methods that incorporate social graphs (G2) consistently outperforms the content-based methods (G1). This signifies the importance of user engagement patterns, and indicates that exploiting social context is central to effective fake news detection. **(3)** Among the social graph based methods (G2), methods that leverage our proposed news engagement graph (GCN, GIN, GraphConv) outperform methods that formulate both news and users as graph nodes (GCNFN and FANG). Our graph formulation is superior, in that it focuses solely on the co-engagement of social users; it facilitates direct information propagation among articles with similar reader groups,




Figure 5: Ablation study of DECOR.

Table 3: Model efficiency comparison on Politifact dataset.

Method	# Params	Runtime (s)	Avg. F1
GCN	49,280	2.47	0.8964
GIN	49,347	3.19	0.9024
GraphConv	98,627	3.72	0.8989
Pro-GNN	296,289	158.03	0.7691
RS-GNN	102,722	39.51	0.7210
GCN w/ DECOR	49,410	4.63	0.9479
GIN w/ DECOR	49,477	5.09	0.9398
GraphConv w/ DECOR	98,757	5.17	0.9377

and avoids the potential task-irrelevant signals from user profiles and related tweets. (4) Existing GSL methods for edge denoising (G3) are not suited to fake news detection. One possible reason is that these methods are similarity-guided, i.e., links between nodes of dissimilar features are strongly suppressed. However, in our fake news detection scenario, two news articles on different topics can be closely connected in terms of co-engagement and veracity type. (5) Compared with competitive fake news detectors, DECOR substantially enhances the performance of three representative GNN backbones. This validates the effectiveness of using degrees and co-engagement to learn a refined news engagement graph.

### 5.3 Ablation Study

We conduct an ablation study to assess the contribution of DECOR’s major components in detecting fake news, and summarize the results in Figure 5. We compare DECOR with two variants, namely **DECOR-COE** without co-engagement, and **DECOR-Deg** without degrees (definitions given in Section 3.3).

As shown in Figure 5, comparing DECOR with either DECOR-COE or DECOR-Deg, the superior fake news detection performance of DECOR illustrates that both co-engagement and degrees play a significant role in achieving the final improvements. Note that in numerous cases, DECOR-COE guided by degrees outperforms the corresponding GNN backbones that utilize the raw news engagement graph, which is consistent with our first empirical finding (Observation 1) on the distinctive connections between node degrees and news veracity. This further highlights the effectiveness of incorporating degree-related properties for social graph refinement.

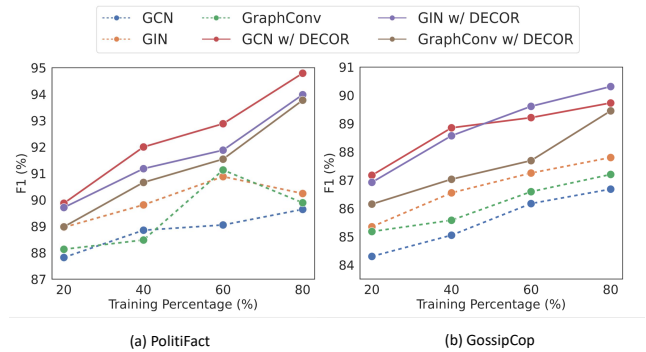


Figure 6: Comparison of DECOR against baselines (F1 Score) under varying training data sizes.

### 5.4 Performance under Label Scarcity

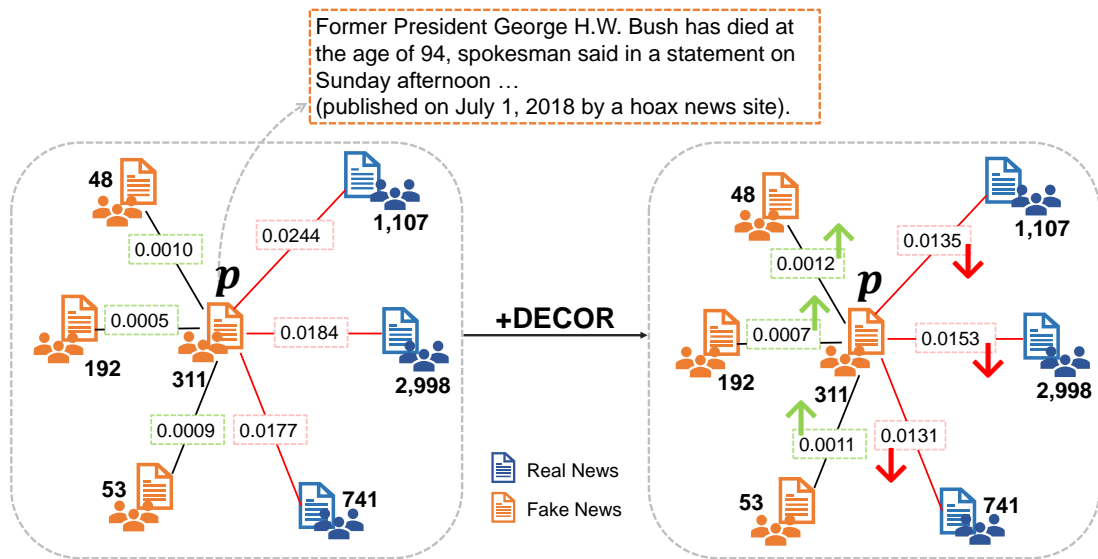
Label scarcity poses an imminent challenge for real-world applications of fake news detection. Due to the timely nature of news articles, high-quality annotations are usually scarce. We evaluate the performance of DECOR under limited training samples, and summarize the results in Figure 6. We observe that DECOR consistently outperforms the competitive GNN baselines on the news engagement graph for all training sizes: 20%, 40%, 60% and 80% of the data. DECOR learns an optimized graph by explicitly leveraging the degree-related structural signals embedded in degrees and news co-engagement, which serves as informative news veracity indicators and thereby complement the limited ground-truth knowledge from fact-checked annotations.

### 5.5 Computational Efficiency

We evaluate the computational cost of DECOR regarding parameter number and model runtime. Specifically, we train all models on the same GPU device for 800 epochs, and compare the time elapsed. Note that both Pro-GNN and RS-GNN adopt the same 2-layer GCN architecture as the “GCN” method reported in Table 3.

Results in Table 3 validate that DECOR is able to achieve impressive performance gains while maintaining low computational cost. Compared with existing GSL methods, three innovations account for DECOR’s efficiency in fake news detection: (1) DECOR leverages low-dimensional features (i.e. degrees and co-engagement) to predict an adjustment score for each edge, whereas existing GSL methods utilize node features that are high-dimensional in terms of news article representations. (2) DECOR utilizes a lightweight degree correction component, which facilitates joint optimization of the social degree correction module and the GNN detector. In contrast, existing GSL methods adopt alternating optimization of the GNN and the link predictor, resulting in slower model training. (3) DECOR operates as pruning on the existing edges in the news engagement graph, whereas existing GSL methods conduct pairwise computations (e.g. feature similarity) among all nodes. Hence, the complexity of DECOR is linear to the number of edges, whereas existing GSL methods incur up to quadratic complexity. These results suggest that DECOR is suitable for deployment in resource-limited scenarios, e.g., online fact-checking services.





**Figure 7: DECOR effectively downweights the noisy edges between influential real and fake news articles, while preserving the informative edges between news of the same veracity type. The edge weights are drawn from the normalized versions of the adjacency matrix  $A$  and the DECOR-refined  $A_c$ , respectively. The number in bold font beside each user icon represents the number of user engagements associated with the corresponding news article.**

## 5.6 Case Study

To further illustrate why DECOR outperforms existing social graph based models and GSL methods, we conduct a case study to illustrate DECOR’s capability of downweighting the noisy edges between fake and real news articles.

In Figure 7, we visualize exemplar cases in the neighborhood of  $p$ , an influential fake news article published by a hoax news site. From the subgraph on the left hand side, we observe that  $p$  is involved in two types of edges: (1) Noisy edges with large edge weights.  $p$  is closely connected with three influential real news pieces. As these articles all focus on trending political topics, they attract a large number of common readers. (2) Clean edges with small edge weights.  $p$  is also connected with several fake news pieces; however, these articles attract less social users, which results in small groups of common readers with  $p$ . These structural patterns are problematic, as propagating information among noisy edges can contaminate the neighborhood, leading to suboptimal article representations. Existing social graph based models generally assume a fixed graph structure and are thereby heavily limited in suppressing edge noise. Prior works on similarity-guided edge denoising also cannot address this issue, as the articles contain similar topics but different veracity. In contrast, DECOR leverages the structural degree-based properties in a flexible manner. This facilitates the elimination of degree-related edge noise. From the subgraph on the right hand side of Figure 7, we find that DECOR effectively suppresses the noisy edges, and recognizes the clean edges by assigning larger weights. These cases provide strong empirical evidence that DECOR effectively refines the news engagement graph for enhanced fake news detection.

## 6 CONCLUSION AND FUTURE WORK

In this paper, we investigate the fake news detection problem from a novel aspect of social graph refinement. We observe that edge noise in the news engagement graph are degree-related, and find that news veracity labels closely correlate with two structural properties: degrees and news co-engagement. Motivated by the DCSBM model’s degree-based probabilistic framework for edge placement, we develop DECOR, a degree-based learnable social graph refinement framework. DECOR facilitates effective suppression of noisy edges through a learnable social degree correction mask, which predicts an adjustment score for each edge based on the aforementioned degree-related properties. Experiments on two real-world benchmarks demonstrate that DECOR can be easily plugged into various powerful GNN backbones as an enhancement. Furthermore, DECOR’s structural corrections are guided by low-dimensional degree-related features, allowing for computationally efficient applications. We believe our empirical and theoretical findings will provide insights for future research in designing and refining more complex multi-relational social graphs for fake news detection.

## 7 ACKNOWLEDGEMENTS

This work was supported by NUS-NCS Joint Laboratory (A-0008542-00-00). The authors would like to thank the anonymous reviewers for their valuable feedback.

## REFERENCES

- [1] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132.
- [2] Yixuan Chen, Dongsheng Li, Peng Zhang, Jie Sui, Qin Lv, Lu Tun, and Li Shang. 2022. Cross-Modal Ambiguity Learning for Multimodal Fake News Detection. In *Proceedings of the ACM Web Conference 2022*. 2897–2905.

- [3] Zhendong Chen, Siu Cheung Hui, Fuzhen Zhuang, Lejian Liao, Fei Li, Meihuizi Jia, and Jiaqi Li. 2022. EvidenceNet: Evidence Fusion Network for Fact Verification. In *Proceedings of the ACM Web Conference 2022*. 2636–2645.
- [4] Niall J. Conroy, Victoria L. Rubin, and Yimin Chen. 2015. Automatic Deception Detection: Methods for Finding Fake News. In *Proceedings of the 78th ASIST Annual Meeting: Information Science with Impact: Research in and for the Community* (St. Louis, Missouri) (ASIST '15). Article 82.
- [5] Limeng Cui, Haeseung Seo, Maryam Tabar, Fenglong Ma, Suhang Wang, and Dongwon Lee. 2020. DETERRENT: Knowledge Guided Graph Attention Network for Detecting Healthcare Misinformation. In *SIGKDD*. 492–502.
- [6] Enyan Dai, Wei Jin, Hui Liu, and Suhang Wang. 2022. Towards Robust Graph Neural Networks for Noisy Graphs with Sparse Labels. In *WSDM*. 181–191.
- [7] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. 2018. Adversarial Attack on Graph Structured Data. In *ICML*. 1115–1124.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*. 4171–4186.
- [9] Yingdong Dou, Kai Shu, Congying Xia, Philip S. Yu, and Lichao Sun. 2021. User Preference-Aware Fake News Detection. In *SIGIR*. 2051–2055.
- [10] Yaqian Dun, Kefei Tu, Chen Chen, Chunyan Hou, and Xiaojie Yuan. 2021. KAN: Knowledge-aware Attention Network for Fake News Detection. In *AAAI*. 81–89.
- [11] Negin Entezari, Saba A. Al-Sayouri, Amirali Darvishzadeh, and Evangelos E. Papalexakis. 2020. All You Need Is Low (Rank): Defending Against Adversarial Attacks on Graphs. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 169–177.
- [12] Kiran Garimella, Gianmarco Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship. *Proceedings of the 2018 World Wide Web Conference*, 913–922.
- [13] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. 2017. Neural Message Passing for Quantum Chemistry. In *ICML*. 1263–1272.
- [14] William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *NeurIPS*. 1025–1035.
- [15] Linmei Hu, Tianchi Yang, Luhao Zhang, Wanjun Zhong, Duyu Tang, Chuan Shi, Nan Duan, and Ming Zhou. 2021. Compare to The Knowledge: Graph Neural Fake News Detection with External Knowledge. In *ACL-IJCNLP*. 754–763.
- [16] Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. 2020. Graph Structure Learning for Robust Graph Neural Networks. In *SIGKDD*. 66–74.
- [17] Brian Karrer and M.E.J. Newman. 2011. Stochastic blockmodels and community structure in networks. *Physical review E, Statistical, nonlinear, and soft matter physics* 83 (01 2011), 016107.
- [18] Diederik Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [19] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- [20] Qifei Li and Wangchunshu Zhou. 2020. Connecting the Dots Between Fact Verification and Fake News Detection. In *Proceedings of the 28th International Conference on Computational Linguistics*. 1820–1825.
- [21] Jasmine McNealy and Michaela Devyn Mullis. 2019. Tea and Turbulence: Communication Privacy Management Theory and Online Celebrity Gossip Forums. *Comput. Hum. Behav.* 92, C (2019), 110–118.
- [22] 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*. 1148–1158.
- [23] Sachin Modgil, Rohit Singh, Shivam Gupta, and Denis Dennehy. 2021. A Confirmation Bias View on Social Media Induced Polarisation During Covid-19. *Information Systems Frontiers* (11 2021).
- [24] Federico Monti, Fabrizio Frasca, Davide Eynard, Damon Mannion, and Michael M. Bronstein. 2019. Fake News Detection on Social Media using Geometric Deep Learning. In *ICLR*.
- [25] Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. 2019. Weisfeiler and Leman Go Neural: Higher-Order Graph Neural Networks. In *AAAI*. 4602–4609.
- [26] 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*. 1165–1174.
- [27] Raymond Nickerson. 1998. Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology* 2 (06 1998), 175–220.
- [28] Kellin Pelrine, Jacob Danovitch, and Reihaneh Rabbany. 2021. The Surprising Performance of Simple Baselines for Misinformation Detection. In *Proceedings of the Web Conference 2021*. 3432–3441.
- [29] Piotr Przybyla. 2020. Capturing the Style of Fake News. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 01 (Apr. 2020), 490–497.
- [30] 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.
- [31] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. arXiv:1910.01108 [cs.CL]
- [32] Lanyu Shang, Ziyi Kou, Yang Zhang, and Dong Wang. 2022. A Duo-Generative Approach to Explainable Multimodal COVID-19 Misinformation Detection. In *Proceedings of the ACM Web Conference 2022*. 3623–3631.
- [33] Qiang Sheng, Juan Cao, Xueyao Zhang, Rundong Li, Danding Wang, and Yongchun Zhu. 2022. Zoom Out and Observe: News Environment Perception for Fake News Detection. In *ACL*. 4543–4556.
- [34] Qiang Sheng, Xueyao Zhang, Juan Cao, and Lei Zhong. 2021. Integrating Pattern and Fact-Based Fake News Detection via Model Preference Learning. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 1640–1650.
- [35] Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019. DEFEND: Explainable Fake News Detection. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 395–405.
- [36] 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 (06 2020), 171–188.
- [37] 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*. 312–320.
- [38] Xianfeng Tang, Yandong Li, Yiwei Sun, Huaxiu Yao, Prasenjit Mitra, and Suhang Wang. 2020. Transferring Robustness for Graph Neural Network Against Poisoning Attacks. In *Proceedings of the 13th International Conference on Web Search and Data Mining*.
- [39] Vaibhav Vaibhav, Raghuram Mandyam, and Eduard Hovy. 2019. Do Sentence Interactions Matter? Leveraging Sentence Level Representations for Fake News Classification. In *Proceedings of the Thirteenth Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-13)*. 134–139.
- [40] Michela Del Vicario, Walter Quattrociocchi, Antonio Scala, and Fabiana Zollo. 2019. Polarization and Fake News: Early Warning of Potential Misinformation Targets. *ACM Trans. Web* 13, 2, Article 10 (2019).
- [41] Binghui Wang and Neil Zhenqiang Gong. 2019. Attacking Graph-Based Classification via Manipulating the Graph Structure. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 2023–2040.
- [42] Yaqing Wang, Fenglong Ma, Haoyu Wang, Kishlay Jha, and Jing Gao. 2021. Multimodal Emergent Fake News Detection via Meta Neural Process Networks. In *SIGKDD*. 3708–3716.
- [43] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. 38–45.
- [44] Huijun Wu, Chen Wang, Yuriy Tyshetskiy, Andrew Docherty, Kai Lu, and Liming Zhu. 2019. Adversarial Examples for Graph Data: Deep Insights into Attack and Defense. In *IJCAI*. 4816–4823.
- [45] Kun Wu, Xu Yuan, and Yue Ning. 2021. Incorporating Relational Knowledge In Explainable Fake News Detection. In *Advances in Knowledge Discovery and Data Mining: 25th Pacific-Asia Conference*. 403–415.
- [46] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *ICLR*.
- [47] Weizhi Xu, Junfei Wu, Qiang Liu, Shu Wu, and Liang Wang. 2022. Evidence-Aware Fake News Detection with Graph Neural Networks. In *Proceedings of the ACM Web Conference 2022*. 2501–2510.
- [48] Ruichao Yang, Xiting Wang, Yiqiao Jin, Chaozhuo Li, Jianxun Lian, and Xing Xie. 2022. Reinforcement Subgraph Reasoning for Fake News Detection. In *SIGKDD*. 2253–2262.
- [49] Shuo Yang, Kai Shu, Suhang Wang, Renjie Gu, Fan Wu, and Huan Liu. 2019. Unsupervised Fake News Detection on Social Media: A Generative Approach. In *AAAI*. 5644–5651.
- [50] Yizhou Zhang, Karishma Sharma, and Yan Liu. 2021. VigDet: Knowledge Informed Neural Temporal Point Process for Coordination Detection on Social Media. In *Advances in Neural Information Processing Systems*.
- [51] Xinyi Zhou, Jindi Wu, and Reza Zafarani. 2020. SAFE: Similarity-Aware Multimodal Fake News Detection. In *Advances in Knowledge Discovery and Data Mining*. 354–367.
- [52] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial Attacks on Neural Networks for Graph Data. In *SIGKDD*. 2847–2856.

**Table 4: Descriptive statistics of news datasets with different topics.**

Dataset	PolitiFact	GossipCop	FANG	MC-Fake (Syria War)	MC-Fake (Health)	MC-Fake (Covid)
# News Articles	497	16,599	727	2,259	5,322	5,248
# Real News	225	12,641	419	2,082	4,788	4,524
# Fake News	272	3,958	308	177	534	724
# User-News Engagements	227,184	963,009	41,747	268,030	632,921	647,387
# Distinct Users	143,481	202,907	30,939	136,847	386,757	303,643

## A EXTENDED ANALYSIS

### A.1 Degree-Related Patterns across News Topics

Recall that we made two degree-related findings in Section 3.3, on how both *degree* and *co-engagement* of news articles closely relate to news veracity. To investigate if these observations are generalizable beyond political news (PolitiFact) and celebrity news (GossipCop), we extend our analysis to four additional news datasets covering three additional topics, namely three datasets from the MC-Fake benchmark [22] on different topics (Syria War, Health and Covid-19) and the FANG dataset [26] (contains news articles about political events and influential rumor events).

As our observations are based on social user engagement patterns, we focus on the news instances with social user engagements, and filter the instances without any user engagement. More specifically, we record social user engagements in terms of source tweets reposting news articles and their retweets, and collect the corresponding user IDs. The descriptive statistics of the datasets are summarized in Table 4.

Following the same procedure of plotting Figure 2 and Figure 3 in Section 3.3, we visualize the node degree distributions of real and fake news via KDE plots in Figure 8, and present the co-engagement patterns of news article pairs in Figure 9. The plots are consistent with our two observations in that (1) the degree distributions of nodes representing fake and real news articles exhibit a clear difference; and (2) given the degrees, edges connecting fake and real news typically have the lowest co-engagement, whereas edges connecting fake news pairs typically have the highest co-engagement. This validates that our observed patterns are widely applicable to news of different topics, and demonstrates promising potential of applying veracity-related co-engagement and degree patterns to refine social graphs that involve news of varying topics.

### A.2 Discussion on Empirical Findings

In this subsection, we discuss the probable reasons leading to our Observation 2 (Section 3.3) on veracity-related patterns between

co-engagement and degrees, which forms the key motivation of DECOR. Recall that Observation 2 is two-fold: (A) Real-Fake news article pairs have the lowest co-engagement given the degrees; and (B) Fake-Fake pairs have higher co-engagement than real-real pairs given the degrees.

We find that both (A) and (B) closely relate with the confirmation bias theory [27], which states that users tend to seek and interpret evidence that upholds their existing beliefs, so as to gain confidence in their biased views.

In terms of (A), as social media platforms foster echo chambers [12] that insulate users from opposing viewpoints, social users tend to repeatedly engage in spreading news articles on certain topics with similar veracity. Hence, social users are less likely to share interest in two news articles of different veracity types (i.e., Real-Fake pairs), which accounts for lower co-engagement than Fake-Fake and Real-Real pairs with the same veracity type.

The underlying phenomena for (B) may vary, as the effects of confirmation bias can be manifested in different forms under different topics. Under topics such as politics [40] or Covid-19 [23], social media platforms can induce opinion polarization, where the user’s attention is highly segregated on a set of certain opinions. In terms of celebrity gossip, certain groups of social media users can engage in boundary coordination to gain control over the information [21]. These phenomena result in sharp community structures, which can be quantified via increased co-engagements.

In conclusion, our second empirical observation can be partially explained by multiple phenomena, and the underlying phenomena can differ across different topics. The benefits of our proposed DECOR framework (Section 4) is that through incorporating a learnable degree correction mechanism, the model is able to recognize the complex veracity-related degree patterns in a more flexible manner. Hence, DECOR facilitates effective detection of news articles without loss of computational efficiency.

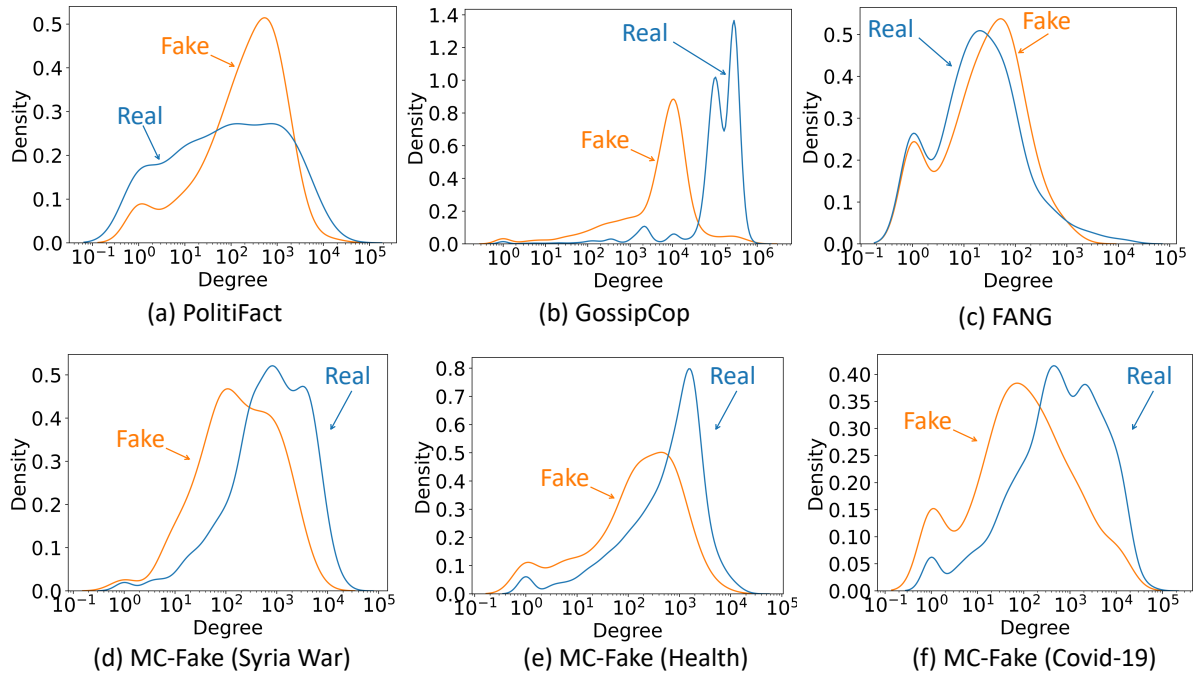


Figure 8: KDE plot of node degree distributions on the news engagement graph.

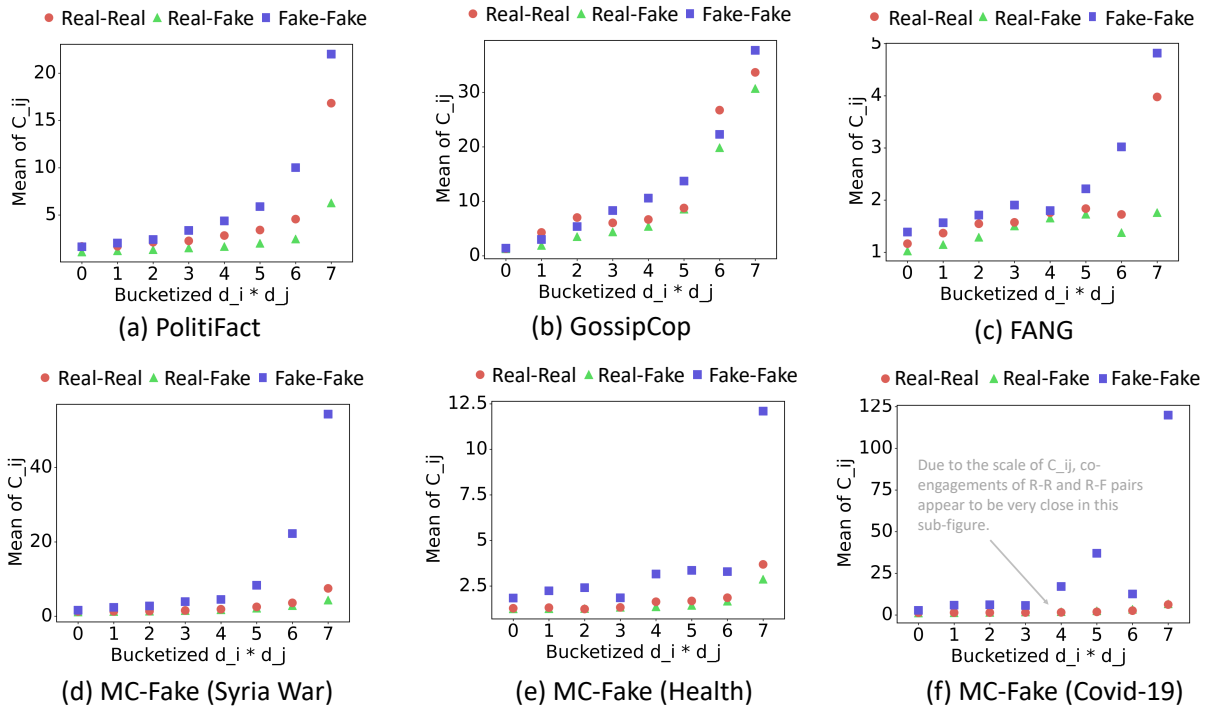


Figure 9: News co-engagement patterns of news article pairs. Edges are grouped based on the connected articles' veracity labels.