

---

# Twibot-22: Towards Graph-Based Twitter Bot Detection

---

Shangbin Feng<sup>1,2\*</sup> Zhaoxuan Tan<sup>1\*</sup> Herun Wan<sup>1\*</sup> Ningnan Wang<sup>1\*</sup> Zilong Chen<sup>1,3\*</sup> Binchi Zhang<sup>1,4\*</sup>  
Qinghua Zheng<sup>1†</sup> Wenqian Zhang<sup>1</sup> Zhenyu Lei<sup>1</sup> Shujie Yang<sup>1</sup> Xinshun Feng<sup>1</sup> Qingyue Zhang<sup>1</sup>  
Hongrui Wang<sup>1</sup> Yuhan Liu<sup>1</sup> Yuyang Bai<sup>1</sup> Heng Wang<sup>1</sup> Zijian Cai<sup>1</sup> Yanbo Wang<sup>1</sup>  
Lijing Zheng<sup>1</sup> Zihan Ma<sup>1</sup> Jundong Li<sup>4</sup> Minnan Luo<sup>1</sup>

Xi'an Jiaotong University<sup>1</sup>, University of Washington<sup>2</sup>, Tsinghua University<sup>3</sup>, University of Virginia<sup>4</sup>  
contact: shangbin@cs.washington.edu

## Abstract

Twitter bot detection has become an increasingly important task to combat misinformation, facilitate social media moderation, and preserve the integrity of the online discourse. State-of-the-art bot detection methods generally leverage the graph structure of the Twitter network, and they exhibit promising performance when confronting novel Twitter bots that traditional methods fail to detect. However, very few of the existing Twitter bot detection datasets are graph-based, and even these few graph-based datasets suffer from limited dataset scale, incomplete graph structure, as well as low annotation quality. In fact, the lack of a large-scale graph-based Twitter bot detection benchmark that addresses these issues has seriously hindered the development and evaluation of novel graph-based bot detection approaches. In this paper, we propose Twibot-22, a comprehensive graph-based Twitter bot detection benchmark that presents the largest dataset to date, provides diversified entities and relations on the Twitter network, and has considerably better annotation quality than existing datasets. In addition, we re-implement 35 representative Twitter bot detection baselines and evaluate them on 9 datasets, including Twibot-22, to promote a fair comparison of model performance and a holistic understanding of research progress. To facilitate further research, we consolidate all implemented codes and datasets into the Twibot-22 evaluation framework, where researchers could consistently evaluate new models and datasets. The Twibot-22 Twitter bot detection benchmark and evaluation framework are publicly available at <https://twibot22.github.io/>.

## 1 Introduction

Automated users on Twitter, also known as Twitter bots, have become a widely known and well-documented phenomenon. Over the past decade, malicious Twitter bots were responsible for a wide range of problems such as online disinformation [Cui et al., 2020, Wang et al., 2020, Lu and Li, 2020], election interference [Howard et al., 2016, Bradshaw et al., 2017, Rossi et al., 2020, Ferrara, 2017], extremism campaign [Ferrara et al., 2016, Marcellino et al., 2020], and even the spread of

---

\* These authors contributed equally to this work.

† Corresponding author: Qinghua Zheng, School of Computer Science and Technology, Xi'an Jiaotong University. Email: qzheng@xjtu.edu.cn; Minnan Luo, School of Computer Science and Technology, Xi'an Jiaotong University. Email: minnluo@xjtu.edu.cn

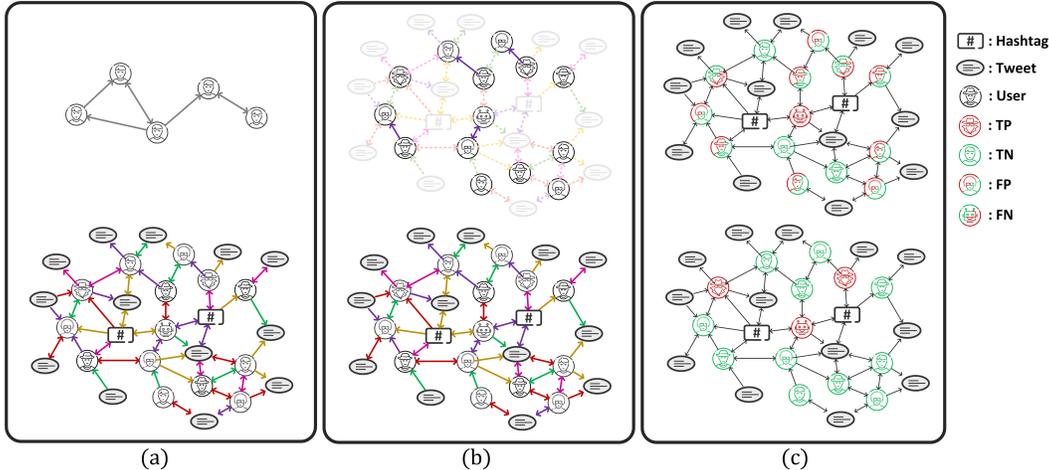


Figure 1: Compared to real-world Twitter (below), existing Twitter bot detection datasets (above) suffer from (a) limited dataset scale, (b) incomplete graph structure, and (c) poor annotation quality.

conspiracy theories [Ferrara, 2020, Ahmed et al., 2020, Anwar et al., 2021]. These societal challenges have called for automatic Twitter bot detection models to mitigate their negative influence.

Existing Twitter bot detection models are generally **feature-based**, where researchers propose to extract numerical features from user information such as metadata [Yang et al., 2020, Lee et al., 2011], user timeline [Mazza et al., 2019, Chavoshi et al., 2016], and follow relationships [Beskow and Carley, 2020, Chu et al., 2012]. However, feature-based approaches are susceptible to adversarial manipulation, *i.e.*, when bot operators try to avoid detection by tampering with these hand-crafted features [Cresci et al., 2017a, Cresci, 2020]. Researchers also proposed **text-based** approaches, where text analysis techniques such as word embeddings [Wei and Nguyen, 2019], recurrent neural networks [Kudugunta and Ferrara, 2018, Feng et al., 2021a], and pre-trained language models [Dukić et al., 2020] are leveraged to analyze tweet content and identify malicious intent. However, new generations of Twitter bots often intersperse malicious content with normal tweets stolen from genuine users [Cresci, 2020, Feng et al., 2021b], thus their bot nature becomes more subtle to text-based methods. With the advent of graph neural networks, recent advances focus on developing **graph-based** Twitter bot detection models. These methods [Ali Alhosseini et al., 2019, Feng et al., 2021b] interpret users as nodes and follow relationships as edges to leverage graph mining techniques such as GCN [Kipf and Welling, 2016], R-GCN [Schlichtkrull et al., 2018], and RGT [Feng et al., 2022] for graph-based bot detection. In fact, recent research have shown that graph-based approaches achieve state-of-the-art performance, are capable of detecting novel Twitter bots, and are better at addressing various challenges facing Twitter bot detection [Feng et al., 2021b, 2022].

However, the development and evaluation of graph-based Twitter bot detection models are poorly supported by existing datasets. The Bot Repository<sup>3</sup> provides a comprehensive collection of existing datasets. Out of the 18 listed datasets, only two of them, TwiBot-20 [Feng et al., 2021c] and cresci-2015 [Cresci et al., 2015], explicitly provide the graph structure among Twitter users. In addition, these two graph-based datasets suffer from the following issues as illustrated in Figure 1:

- **(a) limited dataset scale.** Twibot-20 [Feng et al., 2021c] contains 11,826 labeled users and cresci-15 [Cresci et al., 2015] contains 7,251 labeled users, while online conversations and discussions about heated topics often involve hundreds of thousands of users [Banda et al., 2021].
- **(b) incomplete graph structure.** Real-world Twitter is a heterogeneous information network that contains many types of entities and relations [Feng et al., 2022], while TwiBot-20 and cresci-15 only provide users and follow relationships.
- **(c) low annotation quality.** TwiBot-20 resorted to crowdsourcing for user annotation, while crowdsourcing leads to significant noise [Graells-Garrido and Baeza-Yates, 2022] and is susceptible to the false positive problem [Rauchfleisch and Kaiser, 2020].

<sup>3</sup><https://botometer.osome.iu.edu/bot-repository/datasets.html>

In light of these challenges, we propose TwiBot-22, a graph-based Twitter bot detection benchmark that addresses these issues. Specifically, TwiBot-22 adopts a two-stage controlled expansion to sample the Twitter network, which results in a dataset that is 5 times the size of the largest existing dataset. TwiBot-22 provides 4 types of entities and 14 types of relations in the Twitter network, which provides the first (truly) heterogeneous graph for Twitter bot detection. Finally, TwiBot-22 adopts the weak supervision learning strategy for data annotation which results in significantly improved annotation quality. To compare TwiBot-22 with existing datasets, we re-implement 35 Twitter bot detection baselines and evaluate them on 9 datasets, including TwiBot-22, to provide a holistic view of research progress and highlight the advantages of TwiBot-22. We consolidate all datasets and implemented codes into the TwiBot-22 evaluation framework to facilitate further research. Our main contributions are summarized as follows:

- We propose TwiBot-22, a graph-based Twitter bot detection dataset that establishes the largest benchmark to date, provides diversified entities and relations in the Twitter network, and has considerably improved annotation quality.
- We re-implement and benchmark 35 existing Twitter bot detection models on 9 datasets, including TwiBot-22, to compare different approaches fairly and facilitate a holistic understanding of research progress in Twitter bot detection.
- We present the TwiBot-22 evaluation framework, where researchers could easily reproduce our results, examine existing datasets and methods, infer on unseen Twitter data, and contribute new datasets and models to the framework.

## 2 Related Work

### 2.1 Twitter Bot Detection

Existing Twitter bot detection methods mainly fall into three categories: **feature-based** methods, **text-based** methods, and **graph-based** methods.

**Feature-based** methods conduct feature engineering with user information and apply traditional classification algorithms for bot detection. Various features are extracted from user metadata [Kudugunta and Ferrara, 2018], tweets [Miller et al., 2014], description [Hayawi et al., 2022], temporal patterns [Mazza et al., 2019], and follow relationships [Feng et al., 2021a]. Later research efforts improve the scalability of feature-based approaches [Yang et al., 2020], automatically discover new bots [Chavoshi et al., 2016], and strike the balance between precision and recall [Morstatter et al., 2016]. However, as bot operators are increasingly aware of these hand-crafted features, novel bots often try to tamper with these features to evade detection [Cresci, 2020]. As a result, feature-based methods struggle to keep up with the arms race of bot evolution [Feng et al., 2021a].

**Text-based** methods use techniques in natural language processing to detect Twitter bots based on tweets and user descriptions. Word embeddings [Wei and Nguyen, 2019], recurrent neural networks [Kudugunta and Ferrara, 2018], the attention mechanism [Feng et al., 2021a], and pre-trained language models [Dukić et al., 2020] are adopted to encode tweets for bot detection. Later research combines tweet representations with user features [Cai et al., 2017], learns unsupervised user representations [Feng et al., 2021a], and attempts to address the multi-lingual issue in tweet content [Knauth, 2019]. However, novel bots begin to counter text-based approaches by diluting malicious tweets with content stolen from genuine users [Cresci, 2020]. In addition, Feng et al. [2021b] shows that analyzing tweet content alone might not be robust and accurate for bot detection.

**Graph-based** methods interpret Twitter as graphs and leverage concepts from network science and geometric deep learning for Twitter bot detection. Node centrality [Dehghan et al., 2022], node representation learning [Pham et al., 2022], graph neural networks (GNNs) [Ali Alhosseini et al., 2019], and heterogeneous GNNs [Feng et al., 2021b] are adopted to conduct graph-based Twitter bot detection. Later research try to combine graph-based and text-based methods [Guo et al., 2021a] or propose new GNN architectures to leverage heterogeneities in the Twitter network [Feng et al., 2022]. Graph-based approaches have shown great promise in tackling various challenges facing Twitter bot detection, such as bot communities and bot disguise [Feng et al., 2021b].

The development and evaluation of these models would not be possible without the many valuable Twitter bot detection datasets that were proposed over the past decade. These datasets mainly focus

on politics and elections in the United States [Yang et al., 2020] and European countries [Cresci et al., 2017a]. Cresci-17 [Cresci et al., 2017a] propose the concept of "social spambots" and presents a widely used dataset with more than one type of bots. TwiBot-20 [Feng et al., 2021c] is the latest and most comprehensive Twitter bot detection dataset that addresses the issue of user diversity in previous datasets. However, among 18 datasets presented in the Bot Repository, the go-to place for Twitter bot detection research, only 2 explicitly provide the graph structure of the Twitter network. In addition, these datasets suffer from limited dataset scale, incomplete graph structure, and low annotation quality while increasingly falling short of consistently benchmarking novel graph-based approaches. In light of these challenges, we present TwiBot-22 to alleviate these issues, promote a rethinking of research progress, and facilitate further research in Twitter bot detection.

## 2.2 Graph-based Social Network Analysis

Users in online social networks interact with each other and become part of the network structure, while the network structure is essential in understanding the patterns of social media [Carrington et al., 2005]. With the advent of geometric deep learning, graph neural networks (GNNs) have become increasingly popular in social network analysis research. Qian et al. [2021] propose to model social media with heterogeneous graphs and leverage relational GNNs for illicit drug trafficker detection. Guo et al. [2021b] propose dual graph enhanced embedding neural network to improve graph representation learning and tackle challenges in click-through rate prediction. Graphs and GNNs are also adopted to detect online fraud [Liu et al., 2021, Li et al., 2021, Wang et al., 2021, Mishra et al., 2021, Dou et al., 2020], combat misinformation [Cui et al., 2020, Wang et al., 2020, Lu and Li, 2020, Varlamis et al., 2022, Hu et al., 2021], and improve recommender systems [Ying et al., 2018, Wu et al.]. The task of Twitter bot detection is no exception, where novel and state-of-the-art approaches are increasingly graph-based [Ali Alhosseini et al., 2019, Magelinski et al., 2020, Feng et al., 2021b, 2022, Lei et al., 2022]. In this paper, we propose the TwiBot-22 benchmark to better support the development and evaluation of graph-based Twitter bot detection models.

## 3 TwiBot-22 Dataset

### 3.1 Data Collection

TwiBot-22 aims to present a large-scale and graph-based Twitter bot detection benchmark. To this end, we adopt a two-stage data collection process. We firstly adopt diversity-aware breadth-first search (BFS) to obtain the user network of TwiBot-22. We then collect additional entities and relations on the Twitter network to enrich the heterogeneity of the TwiBot-22 network.

**User network collection.** A common drawback of existing Twitter bot detection datasets is that they only feature a few types of bots and genuine users, while real-world Twitter is home to diversified users and bots [Feng et al., 2021c]. As a result, TwiBot-20 proposes to use breadth-first search (BFS) for user collection, starting from "seed users" and expanding with user follow relationships. To ensure that TwiBot-22 includes different types of bots and genuine users, we augment BFS with two diversity-aware sampling strategies:

- **Distribution diversity.** Given user metadata such as follower count, different types of users fall differently into the metadata distribution. Distribution diversity aims to sample users in the top, middle, and bottom of the distribution. For numerical metadata, among a user’s neighbors and their metadata values, we select the  $k$  users with the highest value,  $k$  with the lowest, and  $k$  randomly chosen from the rest. For true-or-false metadata, we select  $k$  with true and  $k$  with false.
- **Value diversity.** Given a user and its metadata, value diversity is adopted so that neighbors with significantly different metadata values are more likely to be included, ensuring the diversity of collected users. For numerical metadata, among expanding user  $u$ ’s neighbors  $X$  and their metadata values  $x^{num}$ , the probability of user  $x \in X$  being sampled is denoted as  $p(x) \propto |u^{num} - x^{num}|$ . For true-or-false metadata we select  $k$  users from the opposite class.

Based on these sampling strategies, TwiBot-22 conducts diversity-aware BFS starting from @NeurIPSConf. For each neighborhood expansion, one metadata and one of the sampling strategies are randomly adopted. The BFS process stops until the user network contains a desirable amount of Twitter users. More information about the user collection process is presented in Appendix A.2.

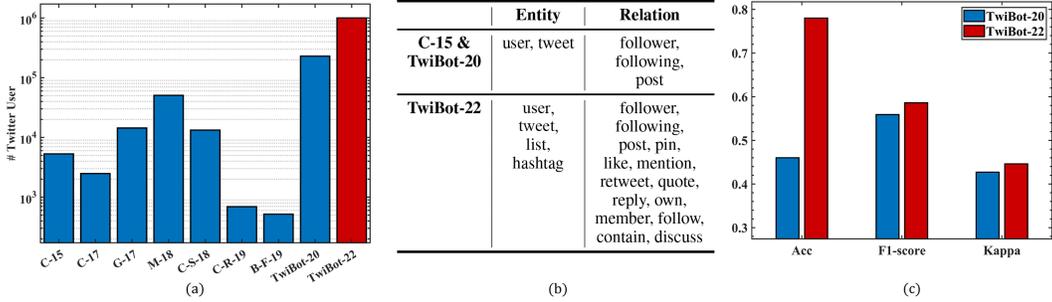


Figure 2: Analyzing TwiBot-22. (a) Dataset scale in terms of users. (b) List of entities and relations in TwiBot-22 and graph-based datasets. (c) Accuracy and F1-score of dataset labels compared to expert annotations as well as Randolph’s kappa coefficient [Randolph, 2005] of expert agreement.

**Heterogeneous graph building.** The user network collection process constructs a homogeneous graph with users as nodes and follow relationships as edges. Apart from that, the Twitter network contains diversified entities and relations such as lists and retweets. Based on the user network, we collect the tweets, associated lists, and mentioned hashtags of these users as well as 12 other relations between users and these new entities. As a result, TwiBot-22 presents a heterogeneous Twitter network with 4 types of entities and 14 types of relations. More information about the heterogeneous Twitter network is presented in Appendix A.4.

As a result, we obtain the TwiBot-22 heterogeneous graph that contains 92,932,326 nodes and 170,185,937 edges. We present more dataset statistics in Table 7 in the appendix.

### 3.2 Data Annotation

Existing bot detection datasets often rely on manual annotation or crowdsourcing, while it is labor-intensive and thus not feasible with the large-scale TwiBot-22. As a result, we adopt weak supervision learning strategy to generate high-quality labels. We firstly invite bot detection experts to annotate 1,000 Twitter users in TwiBot-22. We then generate noisy labels with the help of bot detection models. Finally, we generate high-quality annotations for TwiBot-22 with the Snorkel framework [Ratner et al., 2017].

**Expert annotation.** We randomly select 1,000 users in TwiBot-22 and assign each user to 5 Twitter bot detection experts to identify if it is a bot. We then create an 8:2 split for these expert annotations as training and test sets. More details about these experts are presented in Appendix A.3.

**Generate noisy labels.** We employ 8 hand-crafted labeling functions and 7 competitive feature-based and neural network-based bot detection models to generate noisy labels. For handcrafted labeling functions, we adopt spam keywords in tweets and user descriptions as well as user categorical features such as verified and sensitive tweets. For feature engineering models, we select features based on users’ metadata such as creation time, follower count and name length. We then adopt Adaboost, random forest, and MLP to result in three feature-based classifiers. For neural network-based models, we follow Feng et al. [2021b] to encode user information and employ MLP, GAT [Veličković et al., 2018], GCN [Kipf and Welling, 2016], and R-GCN [Schlichtkrull et al., 2018] as four classifiers. We train these classifiers on the training set of expert annotations and calculate the uncertainty scores for all users in TwiBot-22 under each classifier as  $\phi = -(\hat{y}_0 \log(\hat{y}_0) + \hat{y}_1 \log(\hat{y}_1))$ , where  $\hat{y}_0$  and  $\hat{y}_1$  denote the probability of being genuine users or bots. For each classifier, we then remove model predictions with the top 40% uncertainty scores to alleviate label noises.

**Majority voting.** After obtaining the noisy labels, we evaluate their plausibility with Snorkel [Ratner et al., 2017] and clean the labels at the same time. The output of the Snorkel system are probabilistic labels, thus we use these labels to train an MLP classifier to obtain the final annotations of TwiBot-22. We further evaluate the annotation quality on the test set of expert annotations and

Table 1: Statistics of the 9 datasets. TwiBot-20 contains unlabelled users so that # User  $\neq$  # Human + # Bot. C-15, G-17, C-17, M-18, C-S-18, C-R-19, B-F-19 are short for cresci-2015, gilani-2017, cresci-2017, midterm-18, cresci-stock-2018, cresci-rtbust-2019, botometer-feedback-2019. C-17 contains only "post" edges between users and tweets, which is not a graph-based dataset.

Dataset	C-15	G-17	C-17	M-18	C-S-18	C-R-19	B-F-19	TwiBot-20	TwiBot-22
# Human	1,950	1,394	3,474	8,092	6,174	340	380	5,237	860,057
# Bot	3,351	1,090	10,894	42,446	7,102	353	138	6,589	139,943
# User	5,301	2,484	14,368	50,538	13,276	693	518	229,580	1,000,000
# Tweet	2,827,757	0	6,637,615	0	0	0	0	33,488,192	88,217,457
# Human Tweet	2,631,730	0	2,839,361	0	0	0	0	33,488,192	81,250,102
# Bot Tweet	196,027	0	3,798,254	0	0	0	0	33,488,192	6,967,355
# Edge	7,086,134	0	6,637,615	0	0	0	0	33,716,171	170,185,937

we obtain an 90.5% accuracy. Compared to the 80% accuracy standard in TwiBot-20 [Feng et al., 2021c], TwiBot-22 has considerably improved annotation quality.

### 3.3 Data Analysis

Existing graph-based Twitter bot detection datasets suffer from limited dataset scale, incomplete graph structure, and low annotation quality. As a result, we examine whether TwiBot-22 has adequately addressed these challenges and present our findings in Figure 2.

**Dataset scale.** Figure 2(a) illustrates the number of Twitter users in TwiBot-22 and existing datasets. It is illustrated that TwiBot-22 establishes the largest Twitter bot detection benchmark to date, with approximately 5 times more users than the second-largest TwiBot-20.

**Graph structure.** Figure 2(b) demonstrates that the TwiBot-22 network contains 4 types of entities and 14 types of relations, resulting in significantly enriched graph structure compared to existing graph-based datasets cresci-15 and TwiBot-20 with only 2 entity types and 3 relation types.

**Annotation quality.** TwiBot-20, the largest graph-based Twitter bot detection benchmark to date, leveraged crowdsourcing for data annotation. To improve label quality, TwiBot-22 uses weak supervision learning strategies and leverages 1,000 expert annotations to guide the process. To examine whether TwiBot-22 has improved annotation quality than TwiBot-20, we ask Twitter bot detection experts to participate in an "expert study", where they are asked to evaluate users in TwiBot-20 and TwiBot-22 to examine how often do experts agree with dataset labels. Figure 2(c) illustrates the results, which shows that these experts find TwiBot-22 to provide more consistent, accurate, and trustworthy data annotations. More details about the expert study are available in Appendix A.5.

## 4 Experiments

### 4.1 Experiment Settings

**Datasets.** We evaluate Twitter bot detection models on all 9 datasets in the Bot Repository that contain both bots and genuine users: cresci-2015 [Cresci et al., 2015], gilani-2017 [Gilani et al., 2017], cresci-2017 [Cresci et al., 2017a,b], midterm-18 [Yang et al., 2020], cresci-stock-2018 [Cresci et al., 2018, 2019], cresci-rtbust-2019 [Mazza et al., 2019], botometer-feedback-2019 [Yang et al., 2019], TwiBot-20 [Feng et al., 2021c], and TwiBot-22. We present dataset details in Table 1. We create a 7:2:1 random split as training, validation, and test set to ensure fair comparison.

**Baselines.** We re-implement and evaluate 35 Twitter bot detection baselines SGBot [Yang et al., 2020], Kudugunta and Ferrara [2018], Hayawi et al. [2022], BotHunter [Beskow and Carley, 2018], NameBot [Beskow and Carley, 2019], Abreu et al. [2020], Cresci et al. [2016], Wei and Nguyen [2019], BGSRD [Guo et al., 2021a], RoBERTa [Liu et al., 2019], T5 [Raffel et al., 2020], Eftimion et al. [2018], Kantepe and Ganiz [2017], Miller et al. [2014], Varol et al. [2017], Kouvela et al. [2020], Ferreira Dos Santos et al. [2019], Lee et al. [2011], LOBO [Echeverriñ a et al., 2018],

Table 2: Average bot detection accuracy and standard deviation of five runs of 35 baseline methods on 9 datasets. **Bold** and underline indicate the highest and second highest performance. The F, T, and G in the "Type" column indicates whether a baseline is feature-based, text-based, or graph-based. Cresci *et al.* and Botometer are deterministic methods or APIs without standard deviation. "/" indicates that the dataset does not contain enough user information to support the baseline. "-" indicates that the baseline is not scalable to the largest TwiBot-22 dataset.

Method	Type	C-15	G-17	C-17	M-18	C-S-18	C-R-19	B-F-19	TwiBot-20	TwiBot-22
SGBot	F	77.1 ( $\pm 0.2$ )	<b>78.6</b> ( $\pm 0.8$ )	92.1 ( $\pm 0.3$ )	<u>99.2</u> ( $\pm 0.0$ )	81.3 ( $\pm 0.1$ )	80.9 ( $\pm 1.5$ )	75.5 ( $\pm 1.9$ )	81.6 ( $\pm 0.5$ )	75.1 ( $\pm 0.1$ )
Kudugunta <i>et al.</i>	F	75.3 ( $\pm 0.1$ )	70.0 ( $\pm 1.1$ )	88.3 ( $\pm 0.2$ )	91.0 ( $\pm 0.5$ )	77.5 ( $\pm 0.1$ )	62.9 ( $\pm 0.8$ )	74.0 ( $\pm 4.7$ )	59.6 ( $\pm 0.7$ )	65.9 ( $\pm 0.0$ )
Hayawi <i>et al.</i>	F	84.3 ( $\pm 0.0$ )	52.7 ( $\pm 0.0$ )	90.8 ( $\pm 0.0$ )	84.6 ( $\pm 0.0$ )	50.0 ( $\pm 0.0$ )	51.2 ( $\pm 0.0$ )	<u>77.0</u> ( $\pm 0.0$ )	73.1 ( $\pm 0.0$ )	76.5 ( $\pm 0.0$ )
BotHunter	F	96.5 ( $\pm 1.2$ )	<u>76.4</u> ( $\pm 1.0$ )	88.1 ( $\pm 0.2$ )	<b>99.3</b> ( $\pm 0.0$ )	81.2 ( $\pm 0.2$ )	<u>81.5</u> ( $\pm 1.7$ )	74.7 ( $\pm 1.0$ )	75.2 ( $\pm 0.4$ )	72.8 ( $\pm 0.0$ )
NameBot	F	77.0 ( $\pm 0.0$ )	60.8 ( $\pm 0.0$ )	76.8 ( $\pm 0.0$ )	85.1 ( $\pm 0.0$ )	55.8 ( $\pm 0.0$ )	63.2 ( $\pm 0.0$ )	71.7 ( $\pm 0.0$ )	59.1 ( $\pm 0.1$ )	70.6 ( $\pm 0.0$ )
Abreu <i>et al.</i>	F	75.7 ( $\pm 0.1$ )	74.3 ( $\pm 0.1$ )	92.7 ( $\pm 0.1$ )	96.5 ( $\pm 0.1$ )	75.4 ( $\pm 0.1$ )	80.9 ( $\pm 0.1$ )	<b>77.4</b> ( $\pm 0.1$ )	73.4 ( $\pm 0.1$ )	70.7 ( $\pm 0.1$ )
Cresci <i>et al.</i>	T	37.0	/	33.5	/	/	/	/	47.8	-
Wei <i>et al.</i>	T	96.1 ( $\pm 1.4$ )	/	89.3 ( $\pm 0.7$ )	/	/	/	/	71.3 ( $\pm 1.6$ )	70.2 ( $\pm 1.2$ )
BGSRD	T	87.8 ( $\pm 0.6$ )	48.5 ( $\pm 8.4$ )	75.9 ( $\pm 0.0$ )	82.9 ( $\pm 1.5$ )	50.7 ( $\pm 1.3$ )	50.0 ( $\pm 4.9$ )	59.6 ( $\pm 3.1$ )	66.4 ( $\pm 1.0$ )	71.9 ( $\pm 1.8$ )
RoBERTa	T	97.0 ( $\pm 0.1$ )	/	97.2 ( $\pm 0.0$ )	/	/	/	/	75.5 ( $\pm 0.1$ )	72.1 ( $\pm 0.1$ )
T5	T	92.3 ( $\pm 0.1$ )	/	96.4 ( $\pm 0.0$ )	/	/	/	/	73.5 ( $\pm 0.1$ )	72.1 ( $\pm 0.1$ )
Efthimion <i>et al.</i>	FT	92.5 ( $\pm 0.0$ )	55.5 ( $\pm 0.0$ )	88.0 ( $\pm 0.0$ )	93.4 ( $\pm 0.0$ )	70.8 ( $\pm 0.0$ )	67.6 ( $\pm 0.0$ )	69.8 ( $\pm 0.0$ )	62.8 ( $\pm 0.0$ )	74.1 ( $\pm 0.0$ )
Kantepe <i>et al.</i>	FT	97.5 ( $\pm 1.3$ )	/	98.2 ( $\pm 1.5$ )	/	/	/	/	80.3 ( $\pm 4.3$ )	76.4 ( $\pm 2.4$ )
Miller <i>et al.</i>	FT	75.5 ( $\pm 0.0$ )	51.0 ( $\pm 0.0$ )	77.1 ( $\pm 0.2$ )	83.7 ( $\pm 0.0$ )	52.5 ( $\pm 0.0$ )	54.4 ( $\pm 0.0$ )	<b>77.4</b> ( $\pm 0.0$ )	64.5 ( $\pm 0.4$ )	30.4 ( $\pm 0.1$ )
Varol <i>et al.</i>	FT	93.2 ( $\pm 0.5$ )	/	/	/	/	/	/	78.7 ( $\pm 0.6$ )	73.9 ( $\pm 0.0$ )
Kouvela <i>et al.</i>	FT	97.8 ( $\pm 0.5$ )	74.7 ( $\pm 0.9$ )	<u>98.4</u> ( $\pm 0.1$ )	97.0 ( $\pm 0.1$ )	79.3 ( $\pm 0.3$ )	79.7 ( $\pm 1.2$ )	71.3 ( $\pm 0.9$ )	84.0 ( $\pm 0.4$ )	76.4 ( $\pm 0.0$ )
Santos <i>et al.</i>	FT	70.8 ( $\pm 0.0$ )	51.4 ( $\pm 0.0$ )	73.8 ( $\pm 0.0$ )	86.6 ( $\pm 0.0$ )	62.5 ( $\pm 0.0$ )	73.5 ( $\pm 0.0$ )	71.7 ( $\pm 0.0$ )	58.7 ( $\pm 0.0$ )	-
Lee <i>et al.</i>	FT	<u>98.2</u> ( $\pm 0.1$ )	74.8 ( $\pm 1.2$ )	<b>98.8</b> ( $\pm 0.1$ )	96.4 ( $\pm 0.1$ )	<u>81.5</u> ( $\pm 0.4$ )	<b>83.5</b> ( $\pm 1.9$ )	75.5 ( $\pm 1.3$ )	77.4 ( $\pm 0.5$ )	76.3 ( $\pm 0.1$ )
LOBO	FT	<b>98.4</b> ( $\pm 0.3$ )	/	96.6 ( $\pm 0.3$ )	/	/	/	/	77.4 ( $\pm 0.2$ )	75.7 ( $\pm 0.1$ )
Moghaddam <i>et al.</i>	FG	73.6 ( $\pm 0.2$ )	/	/	/	/	/	/	74.0 ( $\pm 0.8$ )	73.8 ( $\pm 0.0$ )
Alhosseini <i>et al.</i>	FG	89.6 ( $\pm 0.6$ )	/	/	/	/	/	/	59.9 ( $\pm 0.6$ )	47.7 ( $\pm 8.7$ )
Knauth <i>et al.</i>	FTG	85.9 ( $\pm 0.0$ )	49.6 ( $\pm 0.0$ )	90.2 ( $\pm 0.0$ )	83.9 ( $\pm 0.0$ )	<b>88.7</b> ( $\pm 0.0$ )	50.0 ( $\pm 0.0$ )	76.0 ( $\pm 0.0$ )	81.9 ( $\pm 0.0$ )	71.3 ( $\pm 0.0$ )
FriendBot	FTG	96.9 ( $\pm 1.1$ )	/	78.0 ( $\pm 1.0$ )	/	/	/	/	75.9 ( $\pm 0.5$ )	-
SATAR	FTG	93.4 ( $\pm 0.5$ )	/	/	/	/	/	/	84.0 ( $\pm 0.8$ )	-
Botometer	FTG	57.9	71.6	94.2	89.5	72.6	69.2	50.0	53.1	49.9
Rodríguez-Ruiz <i>et al.</i>	FTG	82.4 ( $\pm 0.0$ )	/	76.4 ( $\pm 0.0$ )	/	/	/	/	66.0 ( $\pm 0.1$ )	49.4 ( $\pm 0.0$ )
GraphHist	FTG	77.4 ( $\pm 0.2$ )	/	/	/	/	/	/	51.3 ( $\pm 0.3$ )	-
EvolveBot	FTG	92.2 ( $\pm 1.7$ )	/	/	/	/	/	/	65.8 ( $\pm 0.6$ )	71.1 ( $\pm 0.1$ )
Dehghan <i>et al.</i>	FTG	62.1 ( $\pm 0.0$ )	/	/	/	/	/	/	<u>86.7</u> ( $\pm 0.1$ )	-
GCN	FTG	96.4 ( $\pm 0.0$ )	/	/	/	/	/	/	77.5 ( $\pm 0.0$ )	78.4 ( $\pm 0.0$ )
GAT	FTG	96.9 ( $\pm 0.0$ )	/	/	/	/	/	/	83.3 ( $\pm 0.0$ )	<u>79.5</u> ( $\pm 0.0$ )
HGT	FTG	96.0 ( $\pm 0.3$ )	/	/	/	/	/	/	<b>86.9</b> ( $\pm 0.2$ )	74.9 ( $\pm 0.1$ )
SimpleHGN	FTG	96.7 ( $\pm 0.5$ )	/	/	/	/	/	/	<u>86.7</u> ( $\pm 0.2$ )	76.7 ( $\pm 0.3$ )
BotRGCN	FTG	96.5 ( $\pm 0.7$ )	/	/	/	/	/	/	85.8 ( $\pm 0.7$ )	<b>79.7</b> ( $\pm 0.1$ )
RGT	FTG	97.2 ( $\pm 0.3$ )	/	/	/	/	/	/	86.6 ( $\pm 0.4$ )	76.5 ( $\pm 0.4$ )

Moghaddam and Abbaspour [2022], Ali Alhosseini et al. [2019], Knauth [2019], FriendBot [Beskow and Carley, 2020], SATAR [Feng et al., 2021a], Botometer [Yang et al., 2022], Rodríguez-Ruiz et al. [2020], GraphHist [Magelinski et al., 2020], EvolveBot [Yang et al., 2013], Dehghan et al. [2022], GCN [Kipf and Welling, 2016], GAT [Veličković et al., 2018], HGT [Hu et al., 2020], SimpleHGN [Lv et al., 2021], BotRGCN [Feng et al., 2021b], RGT [Feng et al., 2022]. These methods leverage different aspects of user information and represent different stages of bot detection research. More details about these baseline methods are available in Appendix B.1.

## 4.2 Experiment Results

We re-implement 35 baseline methods and evaluate them on 9 Twitter bot detection datasets. We run each baseline method for **five times** and report the average performance and standard deviation. Table 2 presents the benchmarking results. Our main discoveries are summarized as follows:

- Graph-based approaches are generally more effective than feature-based or text-based methods. As a matter of fact, all top 5 models on TwiBot-20 and TwiBot-22 are graph-based. On average, these top-5 graph-based methods outperform the global average of all baselines by 13.8% and 8.2% on TwiBot-20 and TwiBot-22. This trend suggests that future research in Twitter bot detection should further examine how users and bots interact on Twitter and the heterogeneous graph structure thus formed.

Table 3: Removing the graph-related model component from graph-based methods (w/o G) while comparing to their original versions (Prev.) on TwiBot-20 and TwiBot-22.

Method	TwiBot-20						TwiBot-22					
	Acc			F1			Acc			F1		
	Prev.	w/o G	Diff.	Prev.	w/o G	Diff.	Prev.	w/o G	Diff.	Prev.	w/o G	Diff.
Ali Alhosseini et al. [2019]	59.9	61.8	+1.9	72.1	70.7	-1.4	70.7	66.9	-3.8	5.7	3.5	-2.2
Moghaddam and Abbaspour [2022]	74.0	72.2	-1.8	77.9	75.8	-2.1	73.8	73.3	-0.5	32.1	32.0	-0.1
Knauth [2019]	81.9	81.4	-0.5	85.2	84.9	-0.3	71.3	71.5	+0.2	37.1	11.3	-25.8
EvolveBot [Yang et al., 2013]	65.8	65.1	-0.7	69.7	69.3	-0.4	71.1	71.0	-0.1	14.1	14.0	-0.1
BotRGCN [Feng et al., 2021b]	85.7	82.6	-3.1	87.3	83.8	-3.5	79.7	75.4	-4.3	57.5	41.2	-16.3
RGT [Feng et al., 2022]	86.6	82.6	-4.0	88.0	83.8	-4.2	76.5	75.4	-1.1	42.9	41.2	-1.7

- Most existing datasets do not provide the graph structure of Twitter users to support graph-based approaches, while TwiBot-22 supports all baseline methods and serve as a comprehensive evaluation benchmark. As novel and state-of-the-art models are increasingly graph-based, future Twitter bot detection datasets should provide the graph structure of real-world Twitter.
- TwiBot-22 establishes the largest benchmark while exposing the scalability issues of baseline methods. For example, Dehghan et al. [2022] achieves near-sota performance on TwiBot-20, while failing to scale to TwiBot-22 as our implementation encounters the out-of-memory problem.
- Performance on TwiBot-22 is on average 2.7% lower than on TwiBot-20 across all baseline methods, which demonstrates that Twitter bot detection is still an open problem that calls for further research. This could be attributed to the fact that Twitter bots are constantly evolving to improve their disguise and evade detection, thus bot detection methods should also adapt and evolve.

### 4.3 Removing Graphs from Baselines

Benchmarking results in Table 2 demonstrate that graph-based approaches generally achieve better performance. To examine the role of graphs in graph-based approaches, we remove the graph component in competitive graph-based methods [Ali Alhosseini et al., 2019, Moghaddam and Abbaspour, 2022, Knauth, 2019, Yang et al., 2013, Feng et al., 2021b, 2022] and report model performance in Table 3. It is demonstrated that:

- All baseline methods exhibit performance drops to different extents on two datasets when the graph component is removed. This indicates that the graph-related components in graph-based approaches contribute to bot detection performance.
- For graph neural network-based approaches BotRGCN [Feng et al., 2021b] and RGT [Feng et al., 2022], the performance drop is generally more severe. This suggests that graph neural networks play an important role in boosting model performance and advancing bot detection research.

More details about how graphs are removed from baseline methods are presented in Appendix B.4.

### 4.4 Generalization Study

The challenge of generalization [Yang et al., 2020, Feng et al., 2021a], *i.e.*, whether Twitter bot detection models perform well on unseen data, is essential in ensuring that bot detection research translates to effective social media moderation and real-world impact. To evaluate the generalization ability of existing Twitter bot detection approaches, we identify 10 sub-communities in the TwiBot-22 network. We then use these sub-communities as folds and examine the performance of several representative models when trained on fold  $i$  and evaluated on fold  $j$ . Figure 3 illustrates that:

- **Graph-based methods are better at generalizing to unseen data.** For example, BotRGCN [Feng et al., 2021b] achieves the best avg score among all baseline methods, outperforming the second-highest RGT by 3.66. This suggests that leveraging the network structure of Twitter might be a potential solution to the generalization challenge.
- **Good model performance does not necessarily translate to good generalization ability.** For example, GAT outperforms LOBO by 5.9% and 3.8% on TwiBot-20 and TwiBot-22 respectively

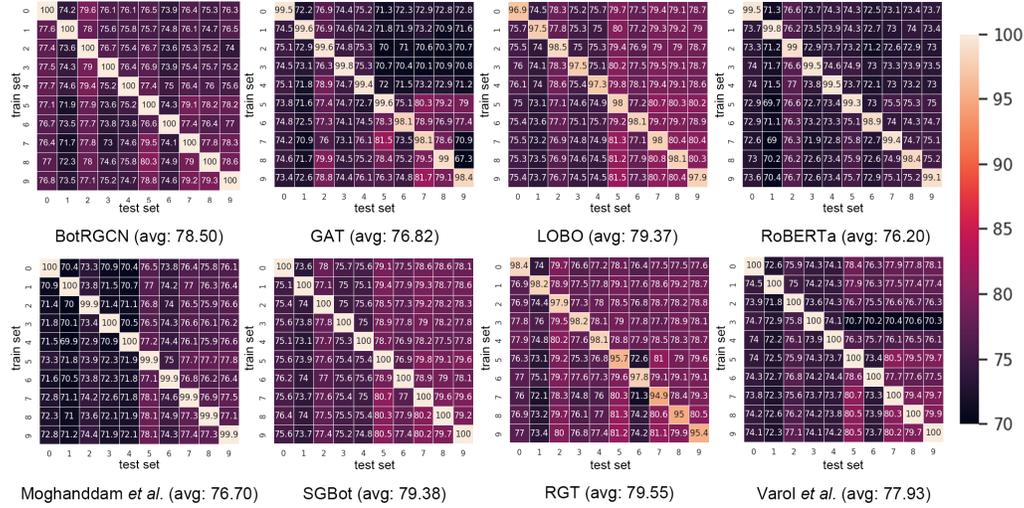


Figure 3: Training models on fold  $i$  and testing on fold  $j$ . We present model accuracy and report the average value of each heatmap (avg), which serves as an overall indicator of generalization ability.

in terms of accuracy. However, GAT has lower avg (-2.55) compared to LOBO. This suggests that future bot detection research should focus on generalization in addition to model performance.

More details about the 10 sub-communities are provided in Appendix B.5.

## 5 Evaluation Framework

We consolidate Twitter bot detection datasets, data preprocessing codes, and all 35 implemented baselines into the TwiBot-22 evaluation framework and make it publicly available. We hope our efforts would facilitate further research in Twitter bot detection through:

- establishing a unified interface for different types of Twitter bot detection datasets
- providing 35 representative baselines and well-documented implementations
- enriching the evaluation framework with new datasets and methods proposed in future research

Please refer to <https://twibot22.github.io/> for more details.

## 6 Conclusion and Future Work

In this paper, we propose TwiBot-22, a graph-based Twitter bot detection benchmark. TwiBot-22 successfully alleviates the challenges of limited dataset scale, incomplete graph structure, and poor annotation quality in existing datasets. Specifically, we employ a two-stage data collection process and adopt the weak supervision learning strategy for data annotation. We then re-implement 35 representative Twitter bot detection models and evaluate them on 9 datasets, including TwiBot-22, to promote a holistic understanding of research progress. We further examine the role of graphs in graph-based methods and the generalization ability of competitive bot detection baselines. Finally, we consolidate all implemented codes into the TwiBot-22 evaluation framework, where researchers could easily reproduce our experiments and quickly test out new datasets and models.

Armed with the TwiBot-22 benchmark and the TwiBot-22 evaluation framework, we aim to investigate these research questions in the future:

- **How do we identify bot clusters and their coordination campaigns?** While existing works study Twitter bot detection through individual analysis, novel Twitter bots are increasingly observed to act in groups and launch coordinated attacks. We aim to complement the scarce literature by proposing temporal and subgraph-level bot detection approaches to address this issue.

- **How do we incorporate multi-modal user features for bot detection?** In addition to text and graph, Twitter users and bots generate multi-modal user information such as images and videos. Since TwiBot-22 provides user media while none of the 35 baselines leverage these modalities, we aim to further explore Twitter bot detection with the help of images and videos.
- **How do we evaluate the generalization ability of bot detection methods?** Existing works mainly focus on bot detection performance while generalization is essential in ensuring that bot detection research generates real-world impact. We aim to complement the scarce literature by proposing measures to quantitatively evaluate bot detection generalization.
- **How do we improve the scalability of graph-based models?** Existing graph-based bot detection methods demand significantly more computation resources and execution time than feature-based models. Given that the Twitter network is rapidly expanding, we aim to further explore scalable and graph-based bot detection methods.

## Acknowledgements

This work was supported by the National Key Research and Development Program of China (No. 2020AAA0108800), National Nature Science Foundation of China (No. 62192781, No. 62272374, No. 61872287, No. 62250009, No. 62137002), Innovative Research Group of the National Natural Science Foundation of China (61721002), Innovation Research Team of Ministry of Education (IRT\_17R86), Project of China Knowledge Center for Engineering Science and Technology and Project of Chinese academy of engineering “The Online and Offline Mixed Educational Service System for The Belt and Road Training in MOOC China”.

We would like to thank the reviewers and area chair for the constructive feedback. We would also like to thank all LUD lab members for our collaborative research environment and for making our bot detection research series possible. As Shangbin concludes his term as the director and Zhaoxuan begins his term, we hope the LUD lab will continue to thrive in the years to come.

## References

- Limeng Cui, Haeseung Seo, Maryam Tabar, Fenglong Ma, Suhang Wang, and Dongwon Lee. Deterrent: Knowledge guided graph attention network for detecting healthcare misinformation. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 492–502, 2020.
- Youze Wang, Shengsheng Qian, Jun Hu, Quan Fang, and Changsheng Xu. Fake news detection via knowledge-driven multimodal graph convolutional networks. In *Proceedings of the 2020 International Conference on Multimedia Retrieval*, pages 540–547, 2020.
- Yi-Ju Lu and Cheng-Te Li. Gcan: Graph-aware co-attention networks for explainable fake news detection on social media. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 505–514, 2020.
- Philip N Howard, Bence Kollanyi, and Samuel Woolley. Bots and automation over twitter during the us election. *Computational propaganda project: Working paper series*, 21(8), 2016.
- Samantha Bradshaw, Bence Kollanyi, Clementine Desigaud, and Gillian Bolsover. Junk news and bots during the french presidential election: What are french voters sharing over twitter? Technical report, COMPROP Data Memo, 2017.
- Sippo Rossi, Matti Rossi, Bikesh Upreti, and Yong Liu. Detecting political bots on twitter during the 2019 finnish parliamentary election. In *Proceedings of the 53rd Hawaii International Conference on System Sciences*, 2020.
- Emilio Ferrara. Disinformation and social bot operations in the run up to the 2017 french presidential election. *arXiv preprint arXiv:1707.00086*, 2017.
- Emilio Ferrara, Wen-Qiang Wang, Onur Varol, Alessandro Flammini, and Aram Galstyan. Predicting online extremism, content adopters, and interaction reciprocity. In *International conference on social informatics*, pages 22–39. Springer, 2016.

- William Marcellino, Madeline Magnuson, Anne Stickells, Benjamin Boudreaux, Todd C Helmus, Edward Geist, and Zev Winkelman. Counter-radicalization bot research using social bots to fight violent extremism. Technical report, Rand Corp Santa Monica CA United States, 2020.
- Emilio Ferrara. What types of covid-19 conspiracies are populated by twitter bots? *arXiv preprint arXiv:2004.09531*, 2020.
- Wasim Ahmed, Francesc López Seguí, Josep Vidal-Alaball, Matthew S Katz, et al. Covid-19 and the film your hospital conspiracy theory: social network analysis of twitter data. *Journal of medical Internet research*, 22(10):e22374, 2020.
- Ahmed Anwar, Haider Ilyas, Ussama Yaqub, and Salma Zaman. Analyzing qanon on twitter in context of us elections 2020: Analysis of user messages and profiles using vader and bert topic modeling. In *DG. O2021: The 22nd Annual International Conference on Digital Government Research*, pages 82–88, 2021.
- Kai-Cheng Yang, Onur Varol, Pik-Mai Hui, and Filippo Menczer. Scalable and generalizable social bot detection through data selection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 1096–1103, 2020.
- Kyumin Lee, Brian Eoff, and James Caverlee. Seven months with the devils: A long-term study of content polluters on twitter. In *Proceedings of the international AAAI conference on web and social media*, volume 5, pages 185–192, 2011.
- Michele Mazza, Stefano Cresci, Marco Avvenuti, Walter Quattrociocchi, and Maurizio Tesconi. Rtbust: Exploiting temporal patterns for botnet detection on twitter. In *Proceedings of the 10th ACM conference on web science*, pages 183–192, 2019.
- Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. Debot: Twitter bot detection via warped correlation. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pages 817–822. IEEE Computer Society, 2016.
- David M Beskow and Kathleen M Carley. You are known by your friends: Leveraging network metrics for bot detection in twitter. In *Open Source Intelligence and Cyber Crime*, pages 53–88. Springer, 2020.
- Zi Chu, Steven Gianvecchio, Haining Wang, and Sushil Jajodia. Detecting automation of twitter accounts: Are you a human, bot, or cyborg? *IEEE Transactions on dependable and secure computing*, 9(6):811–824, 2012.
- Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In *Proceedings of the 26th international conference on world wide web companion*, pages 963–972, 2017a.
- Stefano Cresci. A decade of social bot detection. *Communications of the ACM*, 63(10):72–83, 2020.
- Feng Wei and Uyen Trang Nguyen. Twitter bot detection using bidirectional long short-term memory neural networks and word embeddings. In *2019 First IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, pages 101–109. IEEE, 2019.
- Sneha Kudugunta and Emilio Ferrara. Deep neural networks for bot detection. *Information Sciences*, 467:312–322, 2018.
- Shangbin Feng, Herun Wan, Ningnan Wang, Jundong Li, and Minnan Luo. Satar: A self-supervised approach to twitter account representation learning and its application in bot detection. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 3808–3817, 2021a.
- David Dukić, Dominik Keča, and Dominik Stipić. Are you human? detecting bots on twitter using bert. In *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 631–636. IEEE, 2020.

- Shangbin Feng, Herun Wan, Ningnan Wang, and Minnan Luo. Botrgcn: Twitter bot detection with relational graph convolutional networks. In *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 236–239, 2021b.
- Seyed Ali Alhosseini, Raad Bin Tareaf, Pejman Najafi, and Christoph Meinel. Detect me if you can: Spam bot detection using inductive representation learning. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 148–153, 2019.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pages 593–607. Springer, 2018.
- Shangbin Feng, Zhaoxuan Tan, Rui Li, and Minnan Luo. Heterogeneity-aware twitter bot detection with relational graph transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3977–3985, 2022.
- Shangbin Feng, Herun Wan, Ningnan Wang, Jundong Li, and Minnan Luo. Twibot-20: A comprehensive twitter bot detection benchmark. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 4485–4494, 2021c.
- Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. Fame for sale: Efficient detection of fake twitter followers. *Decision Support Systems*, 80:56–71, 2015.
- Juan M Banda, Ramya Tekumalla, Guanyu Wang, Jingyuan Yu, Tuo Liu, Yuning Ding, Ekaterina Artemova, Elena Tutubalina, and Gerardo Chowell. A large-scale covid-19 twitter chatter dataset for open scientific researchan international collaboration. *Epidemiologia*, 2(3):315–324, 2021.
- Eduardo Graells-Garrido and Ricardo Baeza-Yates. Bots don’t vote, but they surely bother! a study of anomalous accounts in a national referendum. *arXiv preprint arXiv:2203.04135*, 2022.
- Adrian Rauchfleisch and Jonas Kaiser. The false positive problem of automatic bot detection in social science research. *PloS one*, 15(10):e0241045, 2020.
- Zachary Miller, Brian Dickinson, William Deitrick, Wei Hu, and Alex Hai Wang. Twitter spammer detection using data stream clustering. *Information Sciences*, 260:64–73, 2014.
- Kadhim Hayawi, Sujith Mathew, Neethu Venugopal, Mohammad M Masud, and Pin-Han Ho. Deep-robot: a hybrid deep neural network model for social bot detection based on user profile data. *Social Network Analysis and Mining*, 12(1):1–19, 2022.
- Fred Morstatter, Liang Wu, Tahora H Nazer, Kathleen M Carley, and Huan Liu. A new approach to bot detection: striking the balance between precision and recall. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 533–540. IEEE, 2016.
- Chiyu Cai, Linjing Li, and Daniel Zeng. Detecting social bots by jointly modeling deep behavior and content information. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1995–1998, 2017.
- Jürgen Knauth. Language-agnostic twitter-bot detection. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 550–558, 2019.
- Ashkan Dehghan, Kinga Siuta, Agata Skorupka, Akshat Dubey, Andrei Betlen, David Miller, Wei Xu, Bogumil Kaminski, and Pawel Pralat. Detecting bots in social-networks using node and structural embeddings. 2022.
- Phu Pham, Loan TT Nguyen, Bay Vo, and Unil Yun. Bot2vec: A general approach of intra-community oriented representation learning for bot detection in different types of social networks. *Information Systems*, 103:101771, 2022.

- Qinglang Guo, Haiyong Xie, Yangyang Li, Wen Ma, and Chao Zhang. Social bots detection via fusing bert and graph convolutional networks. *Symmetry*, 14(1):30, 2021a.
- Peter J Carrington, John Scott, and Stanley Wasserman. *Models and methods in social network analysis*, volume 28. Cambridge university press, 2005.
- Yiyue Qian, Yiming Zhang, Yanfang Ye, and Chuxu Zhang. Distilling meta knowledge on heterogeneous graph for illicit drug trafficker detection on social media. *Advances in Neural Information Processing Systems*, 34, 2021.
- Wei Guo, Rong Su, Renhao Tan, Huifeng Guo, Yingxue Zhang, Zhirong Liu, Ruiming Tang, and Xiuqiang He. Dual graph enhanced embedding neural network for ctr prediction. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 496–504, 2021b.
- Can Liu, Li Sun, Xiang Ao, Jinghua Feng, Qing He, and Hao Yang. Intention-aware heterogeneous graph attention networks for fraud transactions detection. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3280–3288, 2021.
- Zhao Li, Haishuai Wang, Peng Zhang, Pengrui Hui, Jiaming Huang, Jian Liao, Ji Zhang, and Jiajun Bu. Live-streaming fraud detection: A heterogeneous graph neural network approach. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3670–3678, 2021.
- Li Wang, Peipei Li, Kai Xiong, Jiashu Zhao, and Rui Lin. Modeling heterogeneous graph network on fraud detection: A community-based framework with attention mechanism. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 1959–1968, 2021.
- Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. Modeling users and online communities for abuse detection: A position on ethics and explainability. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3374–3385, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-emnlp.287. URL <https://aclanthology.org/2021.findings-emnlp.287>.
- Yingtong Dou, Zhiwei Liu, Li Sun, Yutong Deng, Hao Peng, and Philip S Yu. Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 315–324, 2020.
- Iraklis Varlamis, Dimitrios Michail, Foteini Glykou, and Panagiotis Tsantilas. A survey on the use of graph convolutional networks for combating fake news. *Future Internet*, 14(3):70, 2022.
- Linmei Hu, Tianchi Yang, Luhao Zhang, Wanjun Zhong, Duyu Tang, Chuan Shi, Nan Duan, and Ming Zhou. Compare to the knowledge: Graph neural fake news detection with external knowledge. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 754–763, 2021.
- Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 974–983, 2018.
- Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. Graph neural networks in recommender systems: a survey. *ACM Computing Surveys (CSUR)*.
- Thomas Magelinski, David Beskow, and Kathleen M Carley. Graph-hist: Graph classification from latent feature histograms with application to bot detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5134–5141, 2020.
- Zhenyu Lei, Herun Wan, Wenqian Zhang, Shangbin Feng, Zilong Chen, Qinghua Zheng, and Minnan Luo. Bic: Twitter bot detection with text-graph interaction and semantic consistency. *arXiv preprint arXiv:2208.08320*, 2022.

- Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. Snorkel: Rapid training data creation with weak supervision. In *Proceedings of the VLDB Endowment. International Conference on Very Large Data Bases*, volume 11, page 269. NIH Public Access, 2017.
- Justus J Randolph. Free-marginal multirater kappa (multirater  $k$  [free]): An alternative to fleiss' fixed-marginal multirater kappa. *Online submission*, 2005.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- Zafar Gilani, Reza Farahbakhsh, Gareth Tyson, Liang Wang, and Jon Crowcroft. Of bots and humans (on twitter). In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, pages 349–354, 2017.
- Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. Social fingerprinting: detection of spambot groups through dna-inspired behavioral modeling. *IEEE Transactions on Dependable and Secure Computing*, 15(4):561–576, 2017b.
- Stefano Cresci, Fabrizio Lillo, Daniele Regoli, Serena Tardelli, and Maurizio Tesconi. Fake: Evidence of spam and bot activity in stock microblogs on twitter. In *Twelfth international AAAI conference on web and social media*, 2018.
- Stefano Cresci, Fabrizio Lillo, Daniele Regoli, Serena Tardelli, and Maurizio Tesconi. Cashtag piggybacking: Uncovering spam and bot activity in stock microblogs on twitter. *ACM Transactions on the Web (TWEB)*, 13(2):1–27, 2019.
- Kai-Cheng Yang, Onur Varol, Clayton A Davis, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Technologies*, 1(1):48–61, 2019.
- David M Beskow and Kathleen M Carley. Bot-hunter: a tiered approach to detecting & characterizing automated activity on twitter. In *Conference paper. SBP-BRiMS: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, volume 3, page 3, 2018.
- David M Beskow and Kathleen M Carley. Its all in a name: detecting and labeling bots by their name. *Computational and Mathematical Organization Theory*, 25(1):24–35, 2019.
- Jefferson Viana Fonseca Abreu, Célia Ghedini Ralha, and João José Costa Gondim. Twitter bot detection with reduced feature set. In *2020 IEEE International Conference on Intelligence and Security Informatics (ISI)*, pages 1–6. IEEE, 2020.
- Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. Dna-inspired online behavioral modeling and its application to spambot detection. *IEEE Intelligent Systems*, 31(5):58–64, 2016.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020.
- Phillip George Efthimion, Scott Payne, and Nicholas Proferes. Supervised machine learning bot detection techniques to identify social twitter bots. *SMU Data Science Review*, 1(2):5, 2018.
- Mücahit Kantepe and Murat Can Ganiz. Preprocessing framework for twitter bot detection. In *2017 International conference on computer science and engineering (ubmk)*, pages 630–634. IEEE, 2017.

- Onur Varol, Emilio Ferrara, Clayton Davis, Filippo Menczer, and Alessandro Flammini. Online human-bot interactions: Detection, estimation, and characterization. In *Proceedings of the international AAAI conference on web and social media*, volume 11, 2017.
- Maria Kouvela, Ilias Dimitriadis, and Athena Vakali. Bot-detective: An explainable twitter bot detection service with crowdsourcing functionalities. In *Proceedings of the 12th International Conference on Management of Digital EcoSystems*, pages 55–63, 2020.
- Eric Ferreira Dos Santos, Danilo Carvalho, Livia Ruback, and Jonice Oliveira. Uncovering social media bots: a transparency-focused approach. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 545–552, 2019.
- Juan Echeverri; a, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Gianluca Stringhini, and Shi Zhou. Lobo: Evaluation of generalization deficiencies in twitter bot classifiers. In *Proceedings of the 34th Annual Computer Security Applications Conference*, pages 137–146, 2018.
- Samaneh Hosseini Moghaddam and Maghsoud Abbaspour. Friendship preference: Scalable and robust category of features for social bot detection. *IEEE Transactions on Dependable and Secure Computing*, 2022.
- Kai-Cheng Yang, Emilio Ferrara, and Filippo Menczer. Botometer 101: Social bot practicum for computational social scientists. *arXiv preprint arXiv:2201.01608*, 2022.
- Jorge Rodríguez-Ruiz, Javier Israel Mata-Sánchez, Raúl Monroy, Octavio Loyola-González, and Armando López-Cuevas. A one-class classification approach for bot detection on twitter. *Computers & Security*, 91:101715, 2020.
- Chao Yang, Robert Harkreader, and Guofei Gu. Empirical evaluation and new design for fighting evolving twitter spammers. *IEEE Transactions on Information Forensics and Security*, 8(8):1280–1293, 2013.
- Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. Heterogeneous graph transformer. In *Proceedings of The Web Conference 2020*, pages 2704–2710, 2020.
- Qingsong Lv, Ming Ding, Qiang Liu, Yuxiang Chen, Wenzheng Feng, Siming He, Chang Zhou, Jianguo Jiang, Yuxiao Dong, and Jie Tang. Are we really making much progress? revisiting, benchmarking and refining heterogeneous graph neural networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1150–1160, 2021.
- Lynnette Hui Xian Ng, Dawn C Robertson, and Kathleen M Carley. Stabilizing a supervised bot detection algorithm: How much data is needed for consistent predictions? *Online Social Networks and Media*, 28:100198, 2022.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf>.
- William Falcon and The PyTorch Lightning team. PyTorch Lightning, 3 2019. URL <https://github.com/PyTorchLightning/pytorch-lightning>.
- Kay Liu, Yingdong Dou, Yue Zhao, Xueying Ding, Xiyang Hu, Ruitong Zhang, Kaize Ding, Canyu Chen, Hao Peng, Kai Shu, et al. Pygod: A python library for graph outlier detection. *arXiv preprint arXiv:2204.12095*, 2022.

- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, October 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6. URL <https://aclanthology.org/2020.emnlp-demos.6>.
- Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Radim Řehůřek and Petr Sojka. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta, May 2010. ELRA. <http://is.muni.cz/publication/884893/en>.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. spaCy: Industrial-strength Natural Language Processing in Python. 2020. doi: 10.5281/zenodo.1212303.
- Joshua Roesslein. tweepy documentation. *Online http://tweepy.readthedocs.io/en/v3.5.2009*.
- Wes McKinney et al. pandas: a foundational python library for data analysis and statistics. *Python for high performance and scientific computing*, 14(9):1–9, 2011.
- Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with numpy. *Nature*, 585(7825):357–362, 2020.
- Clayton Hutto and Eric Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, volume 8, pages 216–225, 2014.
- Guillaume Lemaître, Fernando Nogueira, and Christos K. Aridas. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *Journal of Machine Learning Research*, 18(17):1–5, 2017.