

# MULBOT: Unsupervised Bot Detection Based on Multivariate Time Series

Lorenzo Mannocci<sup>\*†</sup>, Stefano Cresci<sup>\*</sup>, Anna Monreale<sup>†</sup>, Athina Vakali<sup>‡</sup> and Maurizio Tesconi<sup>\*</sup>

<sup>\*</sup> Institute for Informatics and Telematics, National Research Council (IIT-CNR), Pisa, Italy [name.surname@iit.cnr.it]

<sup>†</sup> Department of Computer Science, University of Pisa, Pisa, Italy anna.monreale@unipi.it, lorenzo.mannocci@phd.unipi.it

<sup>‡</sup> School of Informatics, Aristotle University of Thessaloniki, Thessaloniki, Greece avakali@csd.auth.gr

**Abstract**—Online social networks are actively involved in the removal of malicious social bots due to their role in the spread of low quality information. However, most of the existing bot detectors are supervised classifiers incapable of capturing the evolving behavior of sophisticated bots. Here we propose MULBOT, an unsupervised bot detector based on multivariate time series (MTS). For the first time, we exploit multidimensional temporal features extracted from user timelines. We manage the multidimensionality with an LSTM autoencoder, which projects the MTS in a suitable latent space. Then, we perform a clustering step on this encoded representation to identify dense groups of very similar users – a known sign of automation. Finally, we perform a binary classification task achieving f1-score = 0.99, outperforming state-of-the-art methods (f1-score  $\leq$  0.97). Not only does MULBOT achieve excellent results in the binary classification task, but we also demonstrate its strengths in a novel and practically-relevant task: detecting and separating different botnets. In this multi-class classification task we achieve f1-score = 0.96. We conclude by estimating the importance of the different features used in our model and by evaluating MULBOT’s capability to generalize to new unseen bots, thus proposing a solution to the generalization deficiencies of supervised bot detectors.

**Index Terms**—bot detection, multivariate time series, unsupervised learning, social media

## I. INTRODUCTION

Online Social Networks (OSN) have become pervasive in our society and an ever greater source of information for many people. Due to their strong impact, it is paramount to address crucial issues such as disinformation, polarization and hate speech. These concerns, also considered by the European Union<sup>1</sup>, have a wide range of consequences, such as threatening our democracies, polarising debates, and putting the health, security, and environment of all the citizens at risk. In this scenario, social bots play an important role, especially in spreading misinformation [1], [2]. It was estimated that between 9% and 17% of Twitter accounts are bots who contribute on average between 16% and 56% of tweets [3], [4]. On Facebook, bots were estimated to correspond to 11% of all accounts [5]. To make matters worse, many different types of bots exist with different characteristics, behaviors, and aims. For example, fake followers are the simplest kind of bots, exploited to increase the number of followers of a target account, make an account more trustworthy and influential, attracting genuine followers [6]. We can find more complex

behaviors in spambots, whose purpose is spreading (malicious) content, increasing the visibility of some public characters, and promoting a certain company’s product [7].

The wide presence of bots in OSNs brought to the application of many machine learning models, mainly based on supervised approaches. The most widely used supervised bot detector is Botometer<sup>2</sup> [8], which is also extensively used to label ground-truth datasets used to train other bot detectors [9]. However, the dependency of other detectors on Botometer, which has been criticized as inaccurate [10], might propagate existing flaws, instead of fixing them. On the other hand, hand-labeling bot datasets does not appear to be a satisfactory solution either, given the existing biases in human labeling and the results showing that even the most experienced OSN users struggle to spot the latest generation of bots [7].

Moreover, one of the main characteristics of bots is that they *evolve over time*, changing their behavior to evade detection [7]. This is the reason why traditional detection methods based on profile features [8], [11], text content [12]–[14], social relationships or interaction graphs [15]–[17] and posting patterns [18] become soon obsolete and new techniques must be developed. There is a struggle between developers of algorithms for bot detection and increasingly sophisticated bots. The demonstration is that from 2017 onwards, there has been an increase in bot detectors based on adversarial approaches [19] as they are suitable for managing this type of task. Novel and more sophisticated bots do not act anymore individually, but as coordinated groups [5], forming botnets to increase their impact [20], such as the Star Wars botnet [21] or the Bursty botnet [22]. Supervised approaches classify individual users, thus failing to recognize coordinated bots. So, developers of bot detectors moved from supervised learning towards unsupervised models that analyze groups of accounts as a whole, rather than individual accounts, also solving the labeling issue [5]. In this evolving scenario, tabular features extracted from the user profile proved insufficient to identify the new bots [5]. Indeed, considering the whole timeline of posts of the users and the synchronicity between coordinated users was shown to bring a deeper knowledge of the users’ past behavior [9], [23], [24].

Finally, different types of bots (e.g., spammers, fake fol-

<sup>1</sup><https://digital-strategy.ec.europa.eu/en/policies/online-disinformation>

<sup>2</sup><https://botometer.osome.iu.edu>

lowers) have different behavioral characteristics. Therefore supervised learning techniques suffer a decrease in performance when trying to detect previously unseen behaviors (i.e., they do not generalize well), as also demonstrated in [25]. The Leave One Botnet Out (LOBO) test was recently proposed to assess the generalization capabilities of bot detectors [26]. The test involves iteratively re-training bot detectors by excluding a specific botnet at each iteration and by evaluating the effectiveness at detecting bots of the excluded botnet.

On the other hand, most unsupervised approaches rely on textual content analysis, which is now unreliable due to recent improvements in natural language generation [27], [28]. Moreover, the few models based on temporal features, which allow tracking the whole timeline of the users, typically exploit just one dimension at a time [24], discarding relevant information and overlooking much of the bots' behavior.

### A. Contributions

Given the reasons mentioned above, we propose MULBOT, an unsupervised approach for bot detection based on multivariate time series (MTS). The MTS represent temporal information of the tweet's timeline of the user. Then the multidimensional nature is addressed with the use of an LSTM autoencoder [29]. We exploit the autoencoder to compress data, reducing the original dimensionality to a smaller latent space. We choose an LSTM architecture since it is the most appropriate architecture for sequential data [30]. Therefore, the autoencoder encodes the MTS in a suitable latent space, where we apply the appropriate clustering algorithm.

MULBOT makes several contributions with respect to the state-of-the-art, as discussed in the following sections.

*a) Unsupervised and Generalization:* MULBOT is completely unsupervised, overcoming the evolving nature of bots, the emergence of new types of bots and the consequent lack of generalization, typical of supervised approaches. In order to verify if MULBOT is a solution to the lack of generalization of supervised approaches, we perform a test inspired by the LOBO test [26]. We demonstrate that in a real scenario, where new kinds of bots arrive over time, the model is robust and can recognize these new ones.

*b) Multivariate Time Series:* MULBOT is the first method in literature, to the best of our knowledge, based on MTS. MTS allow us to consider multiple dimensions of the users' tweets, such as the number of daily replies, hashtags, URLs and more. Previous works of the state-of-the-art, based on time series, only consider one aspect at a time, losing important information regarding bots' behavior.

*c) Multi-class Classification:* MULBOT is the first unsupervised method that addresses multi-class bot detection. Actually, even within the supervised methods, only [31] and [25] address this task. However, both works solve the task, switching to a multiple binary classification problem. Here instead, we can separate not only bots and genuine users but also recognize different types of bots in a truly multi-class classification task. The latter is a fundamental endeavor because different types of bots have different nature, behaviors,

and aims. Recognizing the kind of bot is a fundamental step in contrasting the activity of malicious bots in OSNs.

*d) Results:* We achieve  $f1\text{-score} = 0.99$  in the traditional binary classification task, overcoming state-of-the-art methods that obtain  $f1\text{-score} \leq 0.97$ . In the new multi-class classification task we achieve  $f1\text{-score} = 0.96$ . In order to make a comparison, we extend the unsupervised method by Ahmed et al. [32] to the multi-class classification task, which obtains  $f1\text{-score} = 0.62$ .

### B. Organization

The rest of this paper is organized as follows. In Section II, we critically discuss the current limits of supervised and unsupervised approaches and the advancement in the state-of-the-art with our approach. In Section III, we present the steps of the proposed method, MULBOT. Section IV describes the performed task, the experimental settings and the implementation of the method presented in Section III. In Section V, we show the reached results for each task. Lately, Section VI concludes the paper discussing possible future works.

## II. RELATED WORKS

In this section, we survey and critically discuss different approaches (mainly unsupervised) used in literature to perform bot detection.

### A. Supervised Approaches

We divide this section into two parts. The first one includes classic supervised machine learning approaches. While the second one presents recent deep learning methods based on neural networks.

*1) Classic Supervised Methods:* There are several works [8], [33]–[36] based on features extraction and classic machine learning classifiers. The most famous supervised approach is Botometer [8], which is often used as ground truth for unsupervised methods. It exploits more than 1,200 features and evaluates users based on their profile information, social network, content, sentiment expressions and the timings of their actions. However, the generality and ease of deployment of this detector are counterbalanced by a reduced bot detection accuracy [7], [37]. In [6], higher performance is reached with a model specialized in a single type of bot, developing a set of supervised machine learning classifiers to detect fake followers, quite easy to identify with respect to other more sophisticated bots. Attempting to address the evolving nature of the bots, in [38] (2013) the authors developed a supervised classifier designed for detecting evolving bots, but yet in 2016, the proposed method was no longer successful at spotting a new wave of malicious accounts [5]. In [31] an attempt to solve the generalization issue is done. Inspired by [25], it trains an ensemble classifier, where each classifier is specialized in a type of bot. This solution only bypasses the problem, reaching a fake generalization. Indeed, a new classifier must be trained for each new type of bot, which is not an optimal solution. However, with the very similar [25], this is the only work where a multi-class classification

is addressed. Actually, in both works, the task is transformed into a multiple binary classification problem through the use of an ensemble classifier, where each internal detector classifies between genuine users and a single type of bot. So, a real multi-class classification is still not addressed.

2) *Neural Network Methods*: In recent years, new approaches based on neural networks have been proposed with the increasing success of deep learning. In [39] a graph-convolutional architecture is used to extract local latent features and classify the graph based on the distribution of these features. Also [40] is based on graph-convolutional networks and [41] is always a graph-based approach, proposing a bot detector that adopts relational graph transformers to leverage the topology and heterogeneity of the real-world Twittersphere. Instead, [42] leverages long short-term memory (LSTM) architecture, exploiting both content and metadata to detect bots at the tweet level. The method tries to minimize the number of features and the size of the training dataset used for classification. Also [43] is based on text and LSTM architecture, while [44] leverages generative adversarial networks. Finally, SATAR [45] is a self-supervised representation learning framework for Twitter users, adapting by pre-training on a massive number of self-supervised users and fine-tuning on detailed bot detection scenarios.

### B. Unsupervised Approaches

The possible unsupervised approaches can be several, based on anomaly detection or graphs that try to extract connectivity patterns of suspicious accounts or even on clustering.

1) *Clustering and Anomaly Detection Methods*: Within works based on clustering, there are [46], [47] and [48]. The latter is based on the fact that spammers do not have high overall peer acceptability and that peers would not accept them in the community. Firstly, they perform the clustering based on user interest distribution, and then they do the spam detection based on peer acceptance. [4] detects tweets containing URLs (likely produced by spambots), with a clustering algorithm to find groups of accounts tweeting texts with high similarity. This method's main issue is using only text features in input. In fact, it has been shown that bot detection based on text features is not a promising approach. Just think of GPT-3 [27], [28], which can generate a text with a so high quality that it can be difficult to determine whether or not it is written by a human.

We can find the same great limit of using text features in [13]. This work proposes a model based on two stream clustering algorithms. They address bot detection as an anomaly detection task treating spambots as outliers, relying on a vector of input features, including content, user information, and tweet text. Among models based on anomaly detection also fall [15] and [49]. The first one is a graph-based approach. The drawback of the second one is that it is not a group-based approach and exploits the behavioral profile of an account, which depends only on its own actions.

Finally, BotWalk [9] is based on the fact that bots are only about 8.5% of Twitter accounts. Therefore, it employs

an ensemble anomaly detection method comparing each user to a seed-set of users (bot and normal users). The exploited features can be divided into 4 categories: metadata-, content-, temporal-, and network-based.

2) *Time Series Based Methods*: We reserve a section for unsupervised methods leveraging temporal features, which are fundamental for detecting novel sophisticated bots. Indeed, the synchronicity of the actions carried on by bots can not be bypassed by bots' developers as there must be a common purpose for the whole coordinated group of bots. So, DeBot [14] works collecting tweets in real-time. It creates a time series for each user, with a one-second sampling rate, where zero value indicates no action and spikes represent the number of actions in that specific second. It does not distinguish different actions but sums up indiscriminately the number of actions in each second. Then it clusters the users hierarchically based on correlation matrix over the set of users. This work is quite essential for our method, MULBOT, because it is the unique case where proper time series are used even if univariate.

Indeed, also in [23] univariate time series are exploited, but with discrete values. It creates a sequence, the so-called DNA-sequence, where each value corresponds to a user's tweet. So the chosen granularity is the tweet, and each sequence is a succession of characters belonging to a predefined alphabet. Each character represents an entity of the tweet (hashtag, URL...). Finally, they apply an unsupervised approach based on longest common substring, which compares different sequences (users). Those who share long behavioral patterns are labeled as spambots, and users who share little similarities are labeled as genuine. If a tweet contains multiple entities, they are managed simply with *tweet contains entities of mixed types*. This kind of choice implies a loss of information regarding the tweet, dealing with a model that only exploits one dimension at a time.

A similar limit can be found in RTbust [24], which constructs a univariate time series based only on retweets. Each value represents the difference between the timestamp of the retweet and a reference timestamp. Then the time series are compressed with an LSTM variational autoencoder exploited to compress the representation and reduce the length of the time series, which all have different lengths, so to obtain a vector of the same length. It uses HDBSCAN [50] to cluster users, classifying noisy points as genuine users and clustered ones as bots.

### C. Advancement of the State-Of-The-Art

Starting from DeBot [14], DNA-sequence [23] and RTbust [24], we can highlight the properties and the deficiencies of the existing time series based model. A summary of the analysis is reported in Table I, where we can observe that none of them considers multiple features at the same time. Indeed, it is true that DeBot considers different properties with a sampling granularity of one second, but it sums up the number of actions without considering their nature. Even DNA-sequence considers different actions, but if there is more

TABLE I  
COMPARISON OF THE METHODS SHARING SOME PROPERTIES WITH MULBOT.

model	approach	time series	multiple features	fine-grained	continuous values	sampling granularity	MTS	multi-class
DeBot [14]	unsupervised	✓	✓	X (sum actions)	X	1 second	X	X
DNA - Sequence [23]	unsupervised	✓	✓	X ("mixed type")	X	tweet	X	X
RTbust [24]	unsupervised	✓	X (retweet)	✓	✓	retweet	X	X
Dimitriadis et al. [31]	supervised	X	✓	✓	-	-	-	✓
MULBOT	unsupervised	✓	✓	✓	✓	1 day	✓	✓

than one action per tweet (such as the presence of a mention and an URL), it assigns a symbol that means "mixed type". Instead, RTbust exploits the timestamp of the retweets, yet not considering other kinds of actions. Therefore, even if these methods initially include multiple features, they actually have a coarse-grained resolution in using features. We can conclude that there is an important lack in the literature, which is the use of temporal information representing the tweets' history of the users considering several features at the same time, such as mentions, hashtags, URLs, replies and more. Finally, it is also included Dimitriadis et al. [31], because it approaches the multi-class classification even if it is a supervised approach based on tabular features.

Therefore, MULBOT, based on temporal information, considers the history of the user. Moreover, these multiple features represented as MTS, capturing complex behaviors, are not aggregated and are considered at a fine-grained level. The unsupervised approach is a solution to the issue of the generalization of the supervised learning method. Finally, to the best of our knowledge, it is the first unsupervised method also designed for the multi-class task, namely the recognition of the kind of bot.

### III. METHOD

We propose MULBOT, an unsupervised approach based on multivariate time series (MTS). As discussed in the earlier sections, the MTS allow taking into account the temporal information of the users, considering more than one feature at a time. In Figure 1, we summarize the proposed framework. Firstly, we extract from the tweets' timeline of the users several temporal information, aggregating them with daily granularity. This allows us to capture complex behaviors of the users and detect even sophisticated bots. Once extracted the time series, we have to manage the multidimensionality of the data. Hence, we train an autoencoder [29], whose encoder maps the MTS into a latent space, where it is easier to perform the last step, which is the clustering algorithm. The considered latent space can have a vectorial dimensionality or can be univariate time series, leaving unaltered the temporal dimension. There is also an optional step between the dimensionality reduction of the autoencoder and the clustering algorithm. In the case of mapping toward univariate time series, we can extract global statistical features and perform the clustering on this final data. In the following sections, we describe each step of the method in detail.

#### A. Step 1: Multivariate Time Series Extraction

The first phase is the retrieval of information from the history of each user's tweets. We extract a set of multivariate time series describing the evolution of specific features that characterize the online user behavior. For each user, we propose to extract temporal features describing tweets, such as the number of URLs or hashtags in tweets aggregated by a daily granularity. The chosen features have already been used successfully in other forms in DNA-sequence [23] or Botwalk [9]. Let  $N$  be the number of users,  $T$  the number of timestamps of the time series and  $D$  the number of input features, the final data shaped as multivariate time series are described by a matrix  $N \times T \times D$ . Being the chosen granularity the day,  $T$  is the number of considered days. For each user  $n = 1 \dots N$ , for each timestamp  $t = 1 \dots T$ , for each input dimension (i.e., feature)  $d = 1 \dots D$ , we define the MTS value as:

$$F_{t,d}^n = \begin{cases} \sum_{j=1}^{TW_t^n} f_{t,d,j}^n, & \text{if } TW_t^n \neq 0 \\ -1, & \text{if } TW_t^n = 0 \end{cases} \quad (1)$$

where  $f_{t,d,j}^n$  is the number of occurrences of the feature  $d$  (e.g. number of hashtags) in the tweet  $j$ , for the user  $n$ , in the day  $t$ .  $TW_t^n$  is the number of tweets of the user  $n$  in the day  $t$ . When  $TW_t^n \neq 0$  is the case in which there is at least a tweet in the day  $t$ , for the user  $n$ . In the case in a certain day, there are no tweets, which is a frequent situation, we assign the special value -1 to each dimension. This setting is useful for distinguishing this special case from the case in which there is a day where the user has tweets but without, for example, URLs, mentions or hashtags. In this last case, a zero value is assigned to the corresponding dimension.

#### B. Step 2: Dimensionality Reduction

Once we extracted the MTS, we have to face the issue of the multidimensional nature of the data. Inspired by RTbust [24], we train an autoencoder to use the encoder to reduce the dimensionality of the time series, mapping the MTS towards a suitable latent space. Since we are dealing with sequential data, the choice for the neural network architecture falls on an LSTM architecture [30]. The first possibility for the dimensionality of the latent space is the univariate time series, represented by a matrix  $N \times T \times 1$ . The second one is the vectorial representation with a chosen latent dimension  $L$ , the

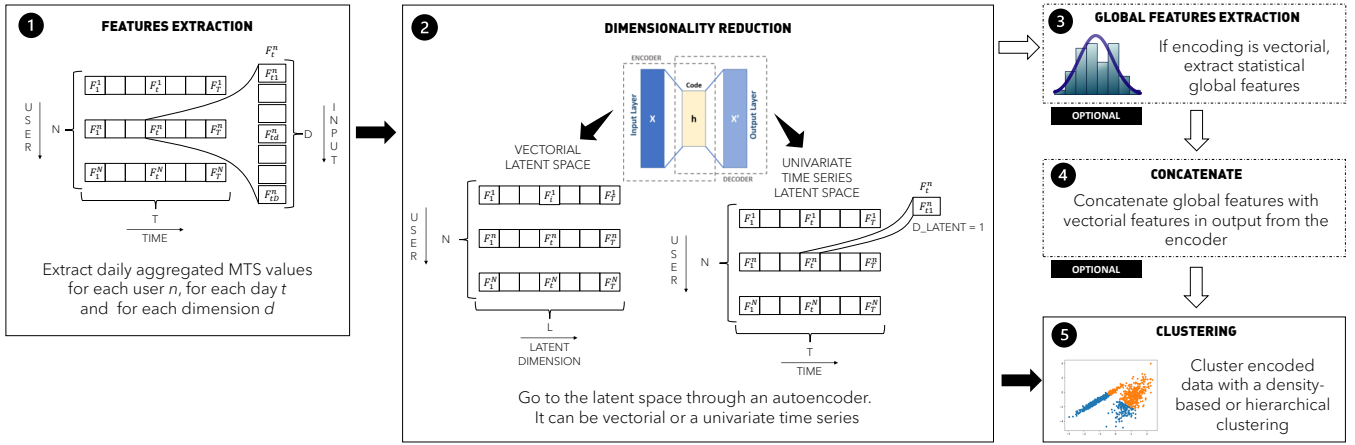


Fig. 1. MULBOT framework, showing the representation of the extracted MTS, the dimensionality reduction step carried on with an autoencoder, the extraction of statistical global features (optional) and finally the clustering algorithm execution.

encoded data have dimension  $N \times L$ . The encoded data is represented in step 2 of Figure 1.

#### C. Step 3-4: Global features and Concatenation (optional)

The third step is an optional phase of the method, and we can execute it only if we map the input data to the univariate time series latent space. Therefore, before the clustering step, we can extract statistical global features<sup>3</sup> from the univariate time series, obtaining tabular/vectorial data. Then, there is a further optional step, the fourth one. We can apply the clustering analysis to this tabular data, or we can concatenate these tabular data with the vectorial data obtained by the encoding of the MTS, using an autoencoder mapping to tabular data.

#### D. Step 5: Clustering Analysis

In the fifth step, data are univariate time series or vectorial features. If we map the data to the vectorial space, we can use a classical clustering algorithm for tabular data. This includes DBSCAN [51], agglomerative hierarchical clustering, HDBSCAN [50], K-Means [52], [53]. HDBSCAN, already used in RTbust [24], is a mix between DBSCAN and hierarchical clustering, which are, for this reason, the most promising ones. These algorithms can also be used in case we map the original data to univariate time series. Indeed, once a similarity measure for MTS is defined, we get a matrix that can be used as input for the above algorithms.

## IV. EXPERIMENTS

In this section, we show the application of MULBOT to a well-known dataset of the literature, the tasks performed, the setting of the hyperparameters and the chosen methods of the state-of-the-art for the comparison of our results.

TABLE II  
SUPPORT AND PERCENTAGE OF GENUINE USERS AND FOR EACH TYPE OF CLASS OF BOTS.

class	bot	type	support	percentage
0	0	Genuine users	3,394	29%
1	1	Spambots 1	991	9%
2		Spambots 2	3,457	30%
3		Spambots 3	464	4%
4		Fake followers	3,202	28%

#### A. Tasks

As discussed in Section III, the MTS are encoded towards a latent space, where it is possible to apply a suitable clustering algorithm. Once the data are clustered, thanks to the use of a labeled dataset, we can evaluate the results with the traditional classification metrics. Therefore, we can perform a binary and a multi-class classification in the final evaluation phase. The binary one is a state-of-the-art task whose aim is the recognition of bots from genuine users. However, here we also address a novel and more difficult task, useful for a deeper knowledge of bots' behaviors. The multi-class classification aims to divide genuine users from bots and also recognize different kind of bots. This is essential in a scenario where each kind of bots behaves in different ways, guided by different purposes [3].

As already mentioned, even if the method is completely unsupervised, with the use of a labeled dataset we can evaluate the results with the traditional metrics used for classification. Therefore, for the binary classification task, we report precision, recall, f1-score and accuracy (Table V). Here, the main considered metric is accuracy, since it is reliable when working with balanced datasets.

<sup>3</sup>An exhaustive list of the extracted features can be found: [https://tsfresh.readthedocs.io/en/latest/api/tsfresh.feature\\_extraction.html#module-tsfresh.feature\\_extraction.feature\\_calculators](https://tsfresh.readthedocs.io/en/latest/api/tsfresh.feature_extraction.html#module-tsfresh.feature_extraction.feature_calculators)

TABLE III  
HYPERPARAMETERS FOR THE AUTOENCODERS WITH UNIVARIATE TIME SERIES AND VECTORIAL LATENT SPACE.

parameter	autoencoder	
	univariate time series	vectorial
activation function	tanh	tanh
output activation function	tanh	tanh
optimizer	rmsprop	rmsprop
learning rate	0.5	0.0002
epochs	250	250
batch size	training set size	training set size
latent dimension	(None, 2976, 1)	300

Instead, in the multi-class classification task (Table VI and Table VII), we also consider the Matthews correlation coefficient (MCC) [54], since our imbalanced dataset would undermine the usefulness of the accuracy metric. As such, for multi-class classification we mainly consider f1-score and MCC.

### B. Dataset

We apply MULBOT on a well-known Twitter dataset in bot detection literature that is *cresci-17*, described in [7]. The *cresci-17* dataset includes genuine users and four kinds of bots: spambots (three types) and fake followers. In Table II, we report the support and the percentage in the dataset of genuine users and each type of bot. Overall, the dataset contains 3,394 genuine users (29% of the total) and 8,144 bots (71%). The large imbalance between bots and genuine users could pose challenges for the binary classification task, reason for which we balance the two classes. Instead, we keep the original (imbalanced) dataset for the multi-class task, since balancing each type of bots would result in discarding too much data.

Therefore, for the binary classification task, to compare our approach with state-of-the-art methods mainly thought for just one kind of bot, we use the *cresci-17* dataset, just including *Spambots 1* and *Genuine users*. Moreover, we balance the dataset with a random downsampling of the genuine users. In this way, we create the best conditions for the state-of-the-art methods for the fairest possible evaluation. Instead, in the multi-class classification task, we use the whole dataset, as we presented here.

### C. Features

The dataset includes several user’s profile features and user tweets’ features, but here we use just the most effective ones according to previous works, such as DNA-sequence [23] or Botwalk [9]. The features initially included in our experiments are: a) *num\_urls*: number of daily URLs b) *num\_hashtags*: number of daily hashtags c) *num\_mentions*: number of daily mentions d) *retweet\_count*: number of daily retweets e) *reply\_count*: number of daily replies f) *favorite\_count*: number of daily favorites.

The chosen granularity is the daily aggregation, considering that a finer granularity would have caused the presence of

TABLE IV  
AUTOENCODERS STRUCTURE IN THE UNIVARIATE TIME SERIES ENCODING AND IN THE VECTORIAL ONE.

layer type	output shape
<i>LSTM Univariate Time Series</i>	
InputLayer	(None, 2976, 6)
LSTM	(None, 2976, 1)
LSTM	(None, 2976, 6)
<i>LSTM Vectorial</i>	
InputLayer	(None, 2976, 6)
LSTM	(None, 2976, 1)
Flatten	(None, 2976)
Dense	(None, 300)
Dense	(None, 2976)
Reshape	(None, 2976, 1)
LSTM	(None, 2976, 3)

many timestamps without tweets and a coarser one would have implied the loss of too much information. Moreover, after several preliminary trials with z-normalization and min-max normalization, the extracted MTS are normalized with min-max normalization, obtaining higher performance.

### D. MULBOT Implementation

In MULBOT, it is required the implementation of the autoencoder and the choice of the suitable clustering algorithm.

The training of the autoencoders [29] is done with a hold-out validation and the mean squared error as loss. The chosen architecture for the neural network is the LSTM, more suitable for time series [30]. Table III shows the chosen values of the hyperparameters for the two autoencoders, one mapping to the univariate time series latent space and the second one to the vectorial encoding. The first column reports the hyperparameters, while the corresponding values for the two autoencoders are reported in the second and the third columns. In table IV, the architecture of the two autoencoders is shown, describing the type of layer and the corresponding output dimension.

Concerning the clustering algorithm, we use for the binary classification task the Agglomerative Hierarchical algorithm with type of *linkage = Ward*<sup>4</sup>. Instead, for the multi-class classification, best results are obtained using a density-based clustering, namely DBSCAN [51]. DBSCAN can cluster data and label specific points as noisy points<sup>5</sup>. From the experiments, we observe that bots are clustered together, showing more similar behaviors with respect to genuine users, who have different behaviors. As a consequence, DBSCAN tends to consider genuine users as noisy points, unable to find a common pattern among them. This result confirms the findings presented in [56], where it is shown that human accounts have more heterogeneous behaviors than automated ones. Given the discovered clusters, the first interpretation we can derive from

<sup>4</sup><https://scikit-learn.org/stable/modules/generated/sklearn.cluster.AgglomerativeClustering.html>

<sup>5</sup>DBSCAN parameters are chosen based on the distribution of distances between points, by following the widely-used method proposed in [55]

TABLE V

RESULTS FOR THE BINARY CLASSIFICATION TASK. IN THE TOP PART, STATE-OF-THE-ART METHODS; IN THE BOTTOM PART, VARIANTS OF MULBOT. BEST RESULTS IN EACH EVALUATION METRIC ARE SHOWN IN BOLD.

methods	type	evaluation metrics			
		precision	recall	f1-score	accuracy
<i>State-of-the-art</i>					
DNA - Sequence [23]	unsupervised	0.982	0.972	0.977	0.976
DNA - Sequence [23]	supervised	0.982	0.977	0.977	0.977
Yang et al. [38]	supervised	0.563	0.170	0.261	0.506
Miller et al. [13]	unsupervised	0.555	0.358	0.435	0.526
Ahmed_MC	unsupervised	0.704	0.699	0.619	0.698
Ahmed_DBSCAN	unsupervised	0.930	0.930	0.930	0.928
<i>MULBOT best variants</i>					
UTS_DBSCAN	unsupervised	0.855	0.855	0.855	0.852
UTS_Hier	unsupervised	0.935	0.930	0.925	0.928
Vec_Hier	unsupervised	0.985	0.985	<b>0.990</b>	0.986
Glob_Hier	unsupervised	<b>0.995</b>	<b>0.995</b>	<b>0.990</b>	<b>0.993</b>
Glob_Vec_Hier	unsupervised	<b>0.995</b>	<b>0.995</b>	<b>0.990</b>	0.992

the results is to consider all noisy points as genuine users and all the others as bots. In this case, we are simply solving a binary classification of each user. The use of a labeled dataset enables a more sophisticated inference about the type of bots represented in any discovered cluster, performing the multi-class classification task. Let  $L = g, b_1, \dots, b_k$  be the set of labels associated with the users, where  $g$  is used to identify genuine users and  $b_i$  represents a particular type of bot. Given the clustering result composed of a set of clusters  $C_0, C_1, C_2, \dots, C_N$ , where  $C_0$  represents the group of points labeled as noisy points and  $C_1, C_2, \dots, C_N$  all the other extracted clusters, we identify the users in  $C_0$  as genuine while the users in each  $C_i$  as bots of type  $b_j$ , if this label is the most frequent in the cluster  $C_i$ .

Lastly, we can optionally extract global statistical features with an automatic extraction process thanks to an existing package<sup>6</sup>, from the encoded univariate time series.

### E. Comparisons

For the binary classification, we compare MULBOT against other methods of the literature, both unsupervised and supervised ones. Namely, we apply the DNA-Sequence [23] both in the unsupervised and supervised approach, Yang et al. [38], Miller et al. [13] and Ahmed et al. [32]. However, this last method expects to use the Markov Clustering, which reaches low performance in all the metrics. Therefore, we replace the Markov Clustering with DBSCAN clustering, increasing the effectiveness of the method.

In the multi-class classification, to the best of our knowledge, there is a lack of methods with which to compare [31]. Therefore, we adapt Ahmed et al. [32], implementing it also for the multi-class classification (always using DBSCAN instead of Markov Clustering). The choice of this method for this task is because it is easily extendable to the multi-class classification task, contrary to most of the methods in literature. We do not implement [25] or [31], both for the high

number of features (more than 1,200 in the first work and about 400 in the second one) and for the type of approach used: multiple binary classification tasks instead of a single multi-class task as in our work.

## V. RESULTS

In this section, we report the results of the tasks and experiments already presented.

### A. Binary Classification

In Table V, we report the results for the binary classification task. In the first part of the table, we present the results of the state-of-the-art methods introduced in Section IV-E. In the second part, we report the best variants of our approach. The following tags for the variants of MULBOT and Ahmed et al. [32] are used:

- UTS\_DBSCAN: Encoding to univariate time series and DBSCAN as clustering algorithm;
- UTS\_Hier: Encoding to univariate time series and Agglomerative Hierarchical clustering algorithm,  $linkage = Ward$ ;
- Vec\_Hier: Encoding to vectorial features and Agglomerative Hierarchical clustering algorithm,  $linkage = Ward$ ;
- Glob\_Hier: Encoding to univariate time series, from which global features are extracted. On these global features it is performed the Agglomerative Hierarchical clustering algorithm,  $linkage = Ward$ ;
- Glob\_Vec\_Hier: Encoding to univariate time series features, from which global features are extracted. These global features are concatenated with the vectorial ones obtained from the corresponding encoding. On this final concatenated data it is performed the Agglomerative Hierarchical clustering algorithm,  $linkage = Ward$ ;
- Ahmed\_MC: Implementation of Ahmed et al. with Markov Clustering according to the original method;
- Ahmed\_DBSCAN: Implementation of Ahmed et al. with DBSCAN to increase the performance.

<sup>6</sup><https://tsfresh.readthedocs.io/en/latest/index.html>

TABLE VI  
WEIGHTED PERFORMANCE OF MULBOT AND AHMED\_DBSCAN FOR MULTI-CLASS CLASSIFICATION.

method	evaluation metrics				
	precision	recall	f1-score	MCC	accuracy
MULBOT	0.966	0.959	0.963	0.949	0.966
Ahmed_DBSCAN	0.704	0.699	0.618	0.647	0.698

TABLE VII  
RESULTS OF MULBOT FOR MULTI-CLASS CLASSIFICATION WITH INPUT FEATURES *retweet\_count*, *reply\_count*, *favorite\_count*, *num\_mentions*, SO EXCLUDING *num\_urls* AND *num\_hashtags*.

class	support	evaluation metrics			
		precision	recall	f1-score	accuracy
0	3,394	0.90	0.99	0.94	0.9624
1	991	0.99	0.85	0.92	
2	3,457	1.00	0.98	0.99	
3	464	1.00	0.82	0.90	
4	3,202	0.99	0.96	0.98	

With Glob\_Hier variant, we reach the highest accuracy of 0.993. The last row of Table V reports the most complex case, in which the global features are concatenated with the vectorial encoding. However, there are no improvements to justify the use of this more complex variant. MULBOT outperforms all the state-of-the-art approaches, both unsupervised and supervised.

### B. Multi-class Classification

This section reports the results of the innovative task of multi-class classification in bot detection. We reach the best performance for the multi-class classification using an LSTM architecture encoding to univariate time series. Moreover the best combination of features in input includes *retweet\_count*, *reply\_count*, *favorite\_count*, *num\_mentions*, so excluding *num\_urls* and *num\_hashtags*. In Table VI the weighted performance of the two methods is reported. We can observe that MULBOT reaches markedly higher results both in accuracy and in f1-score with respect to Ahmed et al. [32]. Even if the used dataset is imbalanced, we can fairly evaluate the results by looking at the MCC (0.949), which is the most robust measure with respect to all the other ones. However, even including all the 6 features presented in Section IV-C, the method reaches 0.95 of f1-score. In Table VII, we report the results of MULBOT for each class, always with the best combination of features in input. For the corresponding class id, refer to Table II.

The method reaches a precision almost equal to one for all the bots. We can better observe it also from the confusion matrix reported in Figure 2. The colors of the cells are based on the normalized confusion matrix on a scale [0, 1]. Instead, the values in the cells are the absolute values of the confusion matrix. Almost all the genuine users are correctly classified (row one of the confusion matrix). This means that the model does not wrongly classify genuine users as bots. This is an

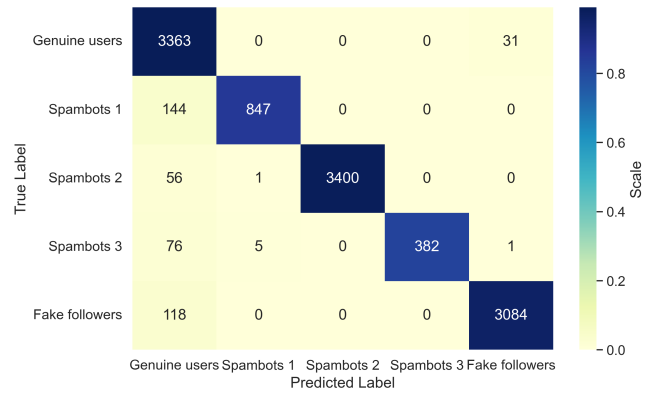


Fig. 2. Confusion matrix of the results of MULBOT for multi-class classification with input features *retweet\_count*, *reply\_count*, *favorite\_count*, *num\_mentions*, so excluding *num\_urls* and *num\_hashtags*. The colors of the cells are based on the normalized confusion matrix, while the values in the cells are the real absolute values.




Fig. 3. Feature importance for the 6 features included in MULBOT.

important result, meaning that in a real scenario, MULBOT is careful not to “ban” genuine Twitter accounts. Instead, bots are more likely to be wrongly classified as genuine users (first column of the confusion matrix). Indeed, precision is 0.90 for *Genuine users* and the recall for *Spambots 1* and *Spambots 3* is respectively 0.85 and 0.82. The method performs well with *Fake followers*, which are one of the most simple existing bots, but also with *Spambots 2*, reaching respectively an f1-score of 0.98 and 0.99.

### C. Feature Importance

This section explores the feature importance of the 6 features included in MULBOT. Inspired by [6], we compute a score for each feature  $i$  as  $S_i = \frac{f_{-i}}{f}$ , where  $f_{-i}$  is the f1-score obtained by MULBOT, excluding the feature  $i$  in input and  $f$  is the f1-score reached by the method with all the 6 features in input. The scores are then normalized and reported in Figure 3. The features *favorite\_count* and *reply\_count* result the most important ones. Instead, *num\_mentions*, *num\_hashtags*, *num\_urls* have almost zero importance. These features are extracted from the tweets’ content, validating that text features are not so important in bot detection tasks.

### D. Generalization of MULBOT: LOBO Test

In this section, we demonstrate how MULBOT can generalize previously unseen bots, overcoming one of the main issues of supervised approaches. With this aim, we perform the LOBO test [26]. In the original work, LOBO test evaluates if



TABLE VIII  
PERCENTAGE CHANGE OF WEIGHTED F1-SCORE FOR EACH LOBO TEST  
INSTANCE (ONE CLASS REMOVED FROM THE TRAINING OF THE  
AUTOENCODER AT A TIME).

excluded class	evaluation metrics	
	accuracy	percentage change
class 1	0.9577	-0.66%
class 2	0.9496	-1.39%
class 3	0.9611	-0.09%
class 4	0.9565	-0.78%

a classifier can detect a bot class, training it only with other bot classes. Inspired by LOBO test, for each kind of bot  $b_i$  in the dataset, we remove the  $b_i$  class before the training of the autoencoder. Then we use the trained encoder to encode also the removed class  $b_i$ . In Table VIII, we report for each experiment (one bot class removed at a time), the percentage change in the weighted f1-score, computed with respect to the MULBOT base case that is the one reported in Table VI. We can observe just small decreases in weighed f1-score. Therefore, our method is robust to the arrival of new types of bots and does not lack generalization, unlike supervised methods. In a real scenario, it means that the unsupervised approach continues working well, even with the arrival on the OSN of new bots with different behaviors. However, we must observe that in case of a high drop in performance, the autoencoder can be easily re-trained. Indeed, thanks to the fact that the whole method is completely unsupervised, there is no need to label the new kinds of bots, which is the main obstacle for the supervised approaches to keep up with the evolutionary nature of the bots.

### E. Results Significance

MULBOT is a group-based approach, being able to recognize new sophisticated bots. The complete unsupervised nature of the method allows reactivity to the evolving nature of the bots. Moreover, the use of MTS is effective since it reaches excellent results both in binary and multi-class scenarios. This highlights how important it is the use of temporal information in the bot detection task, taking into account the whole user’s timeline. The feature importance analysis confirmed the low importance of features extracted from the text, giving more importance to actions performed by other users, such as replies. Including a higher number of features in the model would allow an analysis of the features importance of a wide range of information in the bot detection task. Finally, as anticipated in Section I-A, the developed method can reach high performance even with previously unseen bots. Therefore, it can be a solution to the issue of the generalization, typical of supervised learning approaches.

## VI. CONCLUSION

We proposed an unsupervised method for bot detection to address the issues of the supervised ones. We also exploited temporal features, using multivariate time series for the first

time in bot detection. We exceeded the results of the state-of-the-art methods in the binary classification task, reaching an f1-score of 0.99. We addressed for the first time a multi-classification task, achieving high results in a complex scenario where another method of the literature failed. We conducted several experiments to find the best combination of features in input, evaluating their importance and impact on the performance. Finally, we proposed the LOBO test as proxy for the generalization, showing how the proposed method can generalize to bots never seen before.

In future works, we would like to extend to multi-class classification other state-of-the-art methods to have more comparisons. We used standard features in input, but we can do a more complex features engineering, increasing the features in input and capturing more complex bots’ behaviors. For example, it would be important to insert features such as the time gap between a tweet and the previous one. Moreover, being the method based on clustering, the insertion of new features could imply issues with the curse of dimensionality. Further experiments include different possibilities for the time granularity of the aggregation. The daily granularity is the most balanced in terms of MTS’s length and sparsity of values. However, we could test an hourly or weekly granularity to observe the effects in the results. Finally, we could reach greater certainty of the method’s effectiveness by testing it with different datasets of the literature.

## REFERENCES

- [1] C. Shao, G. L. Ciampaglia, O. Varol, K.-C. Yang, A. Flammini, and F. Menczer, “The spread of low-credibility content by social bots,” *Nature Communications*, vol. 9, no. 1, 2018.
- [2] M. Stella, E. Ferrara, and M. D. Domenico, “Bots increase exposure to negative and inflammatory content in online social systems,” *PNAS*, vol. 115, no. 49, 2018.
- [3] O. Varol, E. Ferrara, C. A. Davis, F. Menczer, and A. Flammini, “Online human-bot interactions: Detection, estimation, and characterization,” in *ICWSM*. AAAI Press, 2017.
- [4] Z. Chen and D. Subramanian, “An unsupervised approach to detect spam campaigns that use botnets on twitter,” 2018.
- [5] S. Cresci, “A decade of social bot detection,” *Communications of the ACM*, vol. 63, no. 10, 2020.
- [6] S. Cresci, R. D. Pietro, M. Petrocchi, A. Spognardi, and M. Tesconi, “Fame for sale: Efficient detection of fake twitter followers,” *Decis. Support Syst.*, vol. 80, 2015.
- [7] —, “The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race,” in *WWW*. ACM, 2017.
- [8] C. A. Davis, O. Varol, E. Ferrara, A. Flammini, and F. Menczer, “BotOrNot: A system to evaluate social bots,” in *WWW (Companion Volume)*. ACM, 2016.
- [9] A. J. Minnich, N. Chavoshi, D. Koutra, and A. Mueen, “Botwalk: Efficient adaptive exploration of Twitter bot networks,” in *ASONAM*. ACM, 2017.
- [10] A. Rauchfleisch and J. Kaiser, “The false positive problem of automatic bot detection in social science research,” *PLoS One*, vol. 15, no. 10, p. e0241045, 2020.
- [11] C. Kater and R. Jäschke, “You shall not pass: detecting malicious users at registration time,” in *Proceedings of the 1st International Workshop on Online Safety, Trust and Fraud Prevention*, 2016.
- [12] S. Lee and J. Kim, “Early filtering of ephemeral malicious accounts on twitter,” *Computer Communications*, vol. 54, 2014.
- [13] Z. Miller, B. Dickinson, W. Deitrick, W. Hu, and A. H. Wang, “Twitter spammer detection using data stream clustering,” *Inf. Sci.*, vol. 260, 2014.
- [14] N. Chavoshi, H. Hamooni, and A. Mueen, “Debot: Twitter bot detection via warped correlation,” in *ICDM*. IEEE Computer Society, 2016.

- [15] M. Jiang, P. Cui, A. Beutel, C. Faloutsos, and S. Yang, "Catching synchronized behaviors in large networks: A graph mining approach," *TKDD*, vol. 10, no. 4, 2016.
- [16] —, "Inferring lockstep behavior from connectivity pattern in large graphs," *KAIS*, vol. 48, no. 2, 2016.
- [17] S. Liu, B. Hooi, and C. Faloutsos, "Holoscope: Topology-and-spike aware fraud detection," in *CIKM*. ACM, 2017.
- [18] G. Stringhini, C. Kruegel, and G. Vigna, "Detecting spammers on social networks," in *ACSAC*, 2010, pp. 1–9.
- [19] S. Cresci, M. Petrocchi, A. Spognardi, and S. Tognazzi, "Better safe than sorry: an adversarial approach to improve social bot detection," in *WebSci*, 2019.
- [20] J. Zhang, R. Zhang, Y. Zhang, and G. Yan, "The rise of social botnets: Attacks and countermeasures," *Trans. Dependable Secur. Comput.*, vol. 15, no. 6, 2018.
- [21] J. Echeverria and S. Zhou, "Discovery, retrieval, and analysis of the 'star wars' botnet in twitter," in *ASONAM*. ACM, 2017.
- [22] J. Echeverria, C. Besel, and S. Zhou, "Discovery of the twitter bursty botnet," in *Data Science for Cyber-Security*. World Scientific, 2019.
- [23] S. Cresci, R. D. Pietro, M. Petrocchi, A. Spognardi, and M. Tesconi, "Social fingerprinting: Detection of spambot groups through DNA-Inspired behavioral modeling," *Trans. Dependable Secur. Comput.*, vol. 15, no. 4, 2018.
- [24] M. Mazza, S. Cresci, M. Avvenuti, W. Quattrociocchi, and M. Tesconi, "RTbust: Exploiting temporal patterns for botnet detection on Twitter," in *WebSci*. ACM, 2019.
- [25] M. Sayyadiharikandeh, O. Varol, K. Yang, A. Flammini, and F. Menczer, "Detection of novel social bots by ensembles of specialized classifiers," in *CIKM*. ACM, 2020.
- [26] J. Echeverría, E. D. Cristofaro, N. Kourtellis, I. Leontiadis, G. Stringhini, and S. Zhou, "LOBO: evaluation of generalization deficiencies in Twitter bot classifiers," in *ACSAC*. ACM, 2018.
- [27] T. Brown *et al.*, "Language models are few-shot learners," in *NEURIPS*, vol. 33. Curran Associates, Inc., 2020.
- [28] T. Fagni, F. Falchi, M. Gambini, A. Martella, and M. Tesconi, "Tweepfake: About detecting deepfake tweets," *PloS one*, vol. 16, no. 5, 2021.
- [29] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, 2006.
- [30] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, 1997.
- [31] I. Dimitriadis, K. Georgiou, and A. Vakali, "Social botomics: A systematic ensemble ml approach for explainable and multi-class bot detection," *Applied Sciences*, vol. 11, no. 21, 2021.
- [32] F. Ahmed and M. Abulaish, "A generic statistical approach for spam detection in online social networks," *Comput. Commun.*, vol. 36, no. 10-11, 2013.
- [33] K.-C. Yang, O. Varol, P.-M. Hui, and F. Menczer, "Scalable and generalizable social bot detection through data selection," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 01, 2020.
- [34] D. M. Beskow and K. M. Carley, "Using random string classification to filter and annotate automated accounts," in *SBP-BRIMS*. Springer, 2018.
- [35] —, "Bot-hunter: a tiered approach to detecting & characterizing automated activity on twitter," in *SBP-BRIMS*, vol. 3, 2018.
- [36] —, "Bot conversations are different: leveraging network metrics for bot detection in twitter," in *ASONAM*. ACM, 2018.
- [37] C. Grimme, D. Assenmacher, and L. Adam, "Changing perspectives: Is it sufficient to detect social bots?" in *SCSM*, vol. 10913. Springer, 2018.
- [38] C. Yang, R. C. Harkreader, and G. Gu, "Empirical evaluation and new design for fighting evolving Twitter spammers," *Trans. Inf. Forensics Secur.*, vol. 8, no. 8, 2013.
- [39] T. Magelinski, D. Beskow, and K. M. Carley, "Graph-hist: Graph classification from latent feature histograms with application to bot detection," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020.
- [40] S. Ali Alhosseini, R. Bin Tareaf, P. Najafi, and C. Meinel, "Detect me if you can: Spam bot detection using inductive representation learning," in *Companion Proceedings of The 2019 World Wide Web Conference*, 2019, pp. 148–153.
- [41] S. Feng, Z. Tan, R. Li, and M. Luo, "Heterogeneity-aware twitter bot detection with relational graph transformers," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 4, 2022.
- [42] S. Kudugunta and E. Ferrara, "Deep neural networks for bot detection," *Information Sciences*, vol. 467, 2018.
- [43] F. Wei and U. T. Nguyen, "Twitter bot detection using bidirectional long short-term memory neural networks and word embeddings," in *TPS-ISA*. IEEE, 2019.
- [44] G. Stanton and A. A. Irissappane, "Gans for semi-supervised opinion spam detection," *arXiv*, 2019.
- [45] S. Feng, H. Wan, N. Wang, J. Li, and M. Luo, "Satar: A self-supervised approach to twitter account representation learning and its application in bot detection," in *CIKM*. ACM, 2021.
- [46] M. C. Benigni, K. Joseph, and K. M. Carley, "Online extremism and the communities that sustain it: Detecting the isis supporting community on twitter," *PloS one*, vol. 12, no. 12, 2017.
- [47] W. Wu, J. Alvarez, C. Liu, and H.-M. Sun, "Bot detection using unsupervised machine learning," *Microsystem Technologies*, vol. 24, no. 1, 2018.
- [48] D. Koggalahewa, Y. Xu, and E. Foo, "An unsupervised method for social network spammer detection based on user information interests," *J. Big Data*, vol. 9, no. 1, 2022.
- [49] X. Ruan, Z. Wu, H. Wang, and S. Jajodia, "Profiling online social behaviors for compromised account detection," *Trans. Inf. Forensics Secur.*, vol. 11, no. 1, 2016.
- [50] L. McInnes, J. Healy, and S. Astels, "HDBSCAN: Hierarchical density based clustering," *J. Open Source Softw.*, vol. 2, no. 11, 2017.
- [51] M. Ester, H. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *KDD*. AAAI Press, 1996.
- [52] J. MacQueen *et al.*, "Some methods for classification and analysis of multivariate observations," in *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, vol. 1, no. 14, 1967.
- [53] S. Lloyd, "Least squares quantization in pcm," *Transactions on information theory*, vol. 28, no. 2, 1982.
- [54] D. Chicco and G. Jurman, "The advantages of the matthews correlation coefficient (mcc) over f1 score and accuracy in binary classification evaluation," *BMC genomics*, vol. 21, no. 1, 2020.
- [55] N. Rahmah and I. S. Sitanggang, "Determination of optimal epsilon (eps) value on dbscan algorithm to clustering data on peatland hotspots in sumatra," in *IOP conference series: earth and environmental science*, vol. 31, no. 1. IOP Publishing, 2016.
- [56] S. Cresci, R. D. Pietro, M. Petrocchi, A. Spognardi, and M. Tesconi, "Emergent properties, models, and laws of behavioral similarities within groups of twitter users," *Comput. Commun.*, vol. 150, 2020.