



Botometer 101: social bot practicum for computational social scientists

Kai-Cheng Yang¹ · Emilio Ferrara² · Filippo Menczer¹

Received: 2 June 2022 / Accepted: 1 August 2022 / Published online: 20 August 2022
© The Author(s), under exclusive licence to Springer Nature Singapore Pte Ltd. 2022

Abstract

Social bots have become an important component of online social media. Deceptive bots, in particular, can manipulate online discussions of important issues ranging from elections to public health, threatening the constructive exchange of information. Their ubiquity makes them an interesting research subject and requires researchers to properly handle them when conducting studies using social media data. Therefore, it is important for researchers to gain access to bot detection tools that are reliable and easy to use. This paper aims to provide an introductory tutorial of Botometer, a public tool for bot detection on Twitter, for readers who are new to this topic and may not be familiar with programming and machine learning. We introduce how Botometer works, the different ways users can access it, and present a case study as a demonstration. Readers can use the case study code as a template for their own research. We also discuss recommended practice for using Botometer.

Keywords Social bots · Twitter · Bot detection · Botometer

Introduction

Social bots are social media accounts controlled in part by software that can post content and interact with other accounts programmatically and possibly automatically [1]. While many social bots are benign, malicious bots can deceptively

✉ Kai-Cheng Yang
yangkc@iu.edu
<https://scholar.google.com/citations?user=tqb96X8AAAAJ&hl=en>

Emilio Ferrara
<https://scholar.google.com/citations?user=0r7Syh0AAAAJ&hl=en>

Filippo Menczer
https://scholar.google.com/citations?user=f_kGJwkAAAAJ&hl=en

¹ Observatory on Social Media, Indiana University Bloomington, Bloomington, IN 47408, USA

² Information Sciences Institute, University of Southern California, Marina del Rey, CA 90292, USA

impersonate humans to manipulate and pollute the information ecosystem. Such malicious bots are involved with all types of online discussions, especially controversial ones. Studies have identified interference of social bots in U.S. elections [2–5], French elections [6], the Brexit referendum [3, 7–9], German elections [10], and the 2017 Catalan referendum [11]. Bots also actively participate in public health debates [12] including those about vaccines [13, 14], the COVID-19 pandemic [15–18], and cannabis [19]. Research has also reported on the presence of social bots in discussions about climate change [20–22], cryptocurrency [23], and the stock market [24, 25].

Malicious social bots demonstrate various behavioral patterns in their actions. They may simply generate a large volume of posts to amplify certain narratives [21, 26] or to manipulate the price of stocks [24, 25] and cryptocurrencies [23]. They can also disseminate low-credibility information strategically by getting involved in the early stage of the spreading process and targeting popular users through mentions and replies [2]. Some bots act as fake followers to inflate the popularity of other accounts [27–29]. In terms of content, malicious bots are found to engage other accounts with negative and inflammatory language [11] or hate speech [17, 30]. In some cases, bots form dense social networks to boost engagement and popularity metrics and to amplify each other's messages [31–33].

Most existing reports and studies on social bots focus on Twitter, largely because its data can be easily accessed. Although Twitter strengthened their efforts to contain malicious actors in recent years,¹ deceptive bots remain prevalent and display evolving tactics to evade detection [34]. This has two implications for researchers. First, characterizing the behavior of and assessing the impact of social bots remains an interesting research topic [35]. Second, researchers need to properly handle bots in their data since their presence may distort analyses [12, 36]. It is therefore crucial for researchers to have access to a reliable tool for detecting social bots.

This practicum aims to provide a tutorial for Botometer, a machine learning tool for bot detection on Twitter. Although other bot detection tools such as tweetbotornot² and Bot Sentinel³ exist, we focus on Botometer for several reasons. First, it is well maintained and has been serving the community for the past seven years without major outages. It has also been routinely upgraded to stay accurate and relevant. Second, Botometer is easily accessible through both a web interface and an application programming interface (API). Anyone with a Twitter account can use the web version for free; researchers with Twitter developer accounts can use the API endpoints to analyze large-scale datasets. The API has a nominal fee for heavy use, which discourages abuse and partially offsets infrastructure and maintenance costs. Third, Botometer is quite popular. It handles around a quarter million daily queries—over half a billion in total since its inception. Finally, Botometer has been extensively validated in the field. Many researchers have applied Botometer in their

¹ <http://blog.twitter.com/common-thread/en/topics/stories/2021/the-secret-world-of-good-bots>.

² An R package for classifying Twitter accounts as bot or not available at github.com/mkearney/Tweetbotornot.

³ A platform that classifies and tracks inauthentic accounts and toxic trolls available at botsentinel.com.

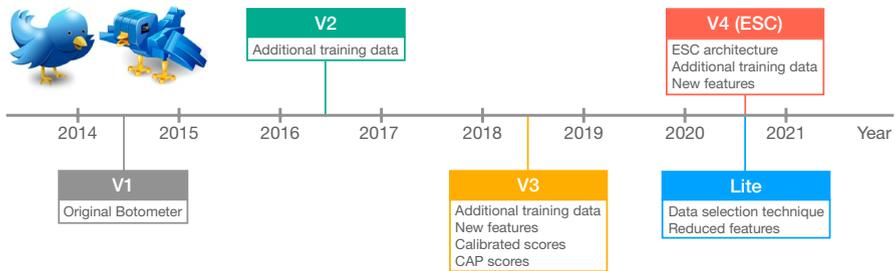


Fig. 1 The timeline of Botometer versions

studies to directly investigate social bots and their impact [10, 13, 19, 25], or to distinguish human accounts and bot-like accounts in order to better address their questions of interest [37–39].

This tutorial is designed for data scientists and computational social scientists who might not be familiar with Botometer, the machine learning methods behind it, its programmatic interface, or how to interpret its results. We start with an introduction to how Botometer works and how users can access it. We then present a case study to demonstrate Botometer usage. The source code for this case study is shared through a public repository for readers to replicate this analysis and use it as a template for their own research. We finally discuss recommended practice.

How Botometer works

Figure 1 presents the timeline and key characteristics of successive Botometer versions over the years. Since the behaviors of bot and human accounts evolve over time, version upgrades are necessary for Botometer to stay accurate and relevant. Upgrades typically included adding new training data and updating model features. The most recent version also involved major architectural changes. Users of Botometer should be aware that results from different versions are usually not comparable and the format of input and output might change as well.

For details of early versions such as V2 [40] and V3 [34], readers can refer to the corresponding papers. This tutorial focuses on V4 [41]. In addition to new training data and new features, this version introduced a new architecture. We will also briefly cover a recently added model for fast bot detection [42].

Supervised machine learning for bot detection

Under the hood, Botometer is a supervised machine learning classifier that distinguishes bot-like and human-like accounts based on their features (i.e., characteristics). Unsupervised methods have also been proposed in the literature [43, 44], but they only allow for the detection of specific, predefined behaviors. Therefore they are not suitable to build a general detection tool.

Table 1 Annotated datasets of human and bot accounts used to train Botometer

Dataset	Bots	Humans	Annotation method	References
varol-icwsm	733	1495	Human annotation	[40]
cresci-17	7049	2764	Various methods	[45]
pronbots	17882	0	Spam bots	[34]
celebrity	0	5918	Celebrity accounts	[34]
vendor-purchased	1087	0	Fake followers	[34]
botometer-feedback	139	380	Human annotation	[34]
political-bots	62	0	Human annotation	[34]
gilani-17	1090	1413	Human annotation	[46]
cresci-rtbust	353	340	Human annotation	[47]
cresci-stock	7102	6174	Signs of coordination	[48]
botwiki	698	0	Self-declared	[42]
midterm-2018	0	7459	Human annotation	[42]
astroturf	505	0	Human annotation	[41]
kaiser	875	499	Politicians + bots	[49]

Botometer considers over 1000 features that can be categorized into six classes: user profile, friends, network, temporal, content and language, and sentiment [40]. For example, the user profile category includes features such as the length of the screen name, whether the account uses the default profile picture and background, the age of the account, etc. The content and language category consists of features such as the number of verbs, nouns, and adjectives in the tweets. For a given account, these features are extracted and encoded as numbers. This way the account can be represented by a vector of feature numbers, enabling machine learning classifiers to process the information.

Supervised machine learning algorithms such as Botometer depend on the availability of training data—accounts labeled as either human or bot. These labels usually come from human annotation [40], automated methods (e.g., honey pots [50]), or botnets that display suspicious behaviors [44, 51]. A critical issue with existing datasets is the lack of ground truth. There is no objective, agreed-upon, operational definition of social bot. A further complicating factor is the prevalence of accounts that lie in the gray area between human and bot behavior, where even experienced researchers cannot easily discriminate. Nevertheless, datasets do include many typical bots; using the training labels as proxies for ground truth makes it possible to build practically viable tools.

Botometer-V4 is trained on a variety of datasets shown in Table 1, which are publicly available in a Bot Repository.⁴ With all training accounts being represented as feature vectors, a classifier can learn the characteristics of bot and human accounts. Botometer uses a classification model called Random Forest, which consists of many rules learned from the training data.

⁴ botometer.osome.iu.edu/bot-repository.

To evaluate a Twitter account, Botometer first fetches its 200 most recent tweets and tweets mentioning it from Twitter, extracts its features from the collected data, and represents this information as a feature vector. Each model rule uses some of the features and provides a vote on whether an account is more similar to bot or human accounts in the training data. Based on how many rules vote for the bot or human class, the model provides a “bot score” between zero and one: a score close to one means the account is highly automated, while a score near zero means a human is likely handling the account. Some accounts may demonstrate the characteristics of both humans and bots. For instance, a bot creator might generate content like a regular user but uses a script to control many accounts. These cases can be confusing for the classifier, which would then produce scores around 0.5.

While human accounts tend to behave similarly, different types of bots usually have unique behavioral patterns. Based on this observation, Botometer-V4 uses several specialized Random Forest classifiers: one for each type of bots in the training data and one for humans. The results of this Ensemble of Specialized Classifiers (ESC) are aggregated to produce a final result. More details about the ESC architecture can be found in the original paper [41]. At the end of the day, the ESC architecture is still a machine learning classifier, which yields scores between 0 and 1. Different from a single Random Forest, the scores generated by ESC tend to have a bimodal distribution.

It is worth mentioning that the content and language features and sentiment features are based on English. When a non-English account is passed to Botometer, these features become meaningless and might affect the classification. As a workaround, Botometer also returns a language-independent score, which is generated without any language-related features. Users need to be aware of the account language and choose the most appropriate Botometer score.

Model accuracy

The accuracy of the model is evaluated through 5-fold cross-validation on the annotated datasets shown in Table 1. Simply speaking, the classifier is trained on part of the annotated datasets and tested on the rest to provide a sense of its accuracy. In the experimental environment, Botometer works really well. V4 has an AUC (area under the receiver operating characteristic curve) of 0.99, suggesting that the model can distinguish bot and human accounts in Table 1—as well as accounts in the wild that resemble those in the training datasets—with very high accuracy.

However, Botometer is not perfect and may misclassify accounts due to several factors. For example, the training datasets might have conflicts because they were created by different people with different standards. In some cases Botometer fails to capture the features that can help distinguishing different accounts. Botometer sometimes struggles with inactive accounts since not enough data is available for evaluation. The accuracy of the model may further decay when dealing with new accounts different from those in the training datasets. These accounts might come from a different context, use different languages other than English [52, 53], or show novel behavioral patterns [34, 45, 54]. These limitations are inevitable for all supervised

machine learning algorithms, and are the reasons why Botometer has to be upgraded routinely.

Some critics exploit these limitations to undermine the entire field of study devoted to social bots. For example, one might select small sets of accounts with large false-positive error rates to argue that no bot detection tool is valid or that social bots do not exist at all. These arguments use fallacies such as cherry-picking and strawman in disingenuous ways. Validation through manual annotations is extremely valuable, especially when highlighting cases where existing machine learning models perform poorly, but should be used in constructive ways. New manually-annotated datasets should be made available, ideally via the public Bot Repository, to support the development of improved models.

Results interpretation

Early versions of Botometer returned to users raw scores in the unit interval, produced by the Random Forest classifiers. Although users often treated them as probabilities, such interpretation is inaccurate. Consider Twitter accounts a and b and their respective scores 0.7 and 0.3 produced by a Random Forest classifier. We can say that a is more bot-like than b , but it is inaccurate to say that there is a 70% chance that a is a bot or that a is 70% bot. Since Botometer-V3, the scores displayed in the web interface are rescaled to the range 0–5 to discourage inaccurate probabilistic interpretations.

For users who need a probabilistic interpretation of a bot score, the Complete Automation Probability (CAP) represents the probability that an account with a given score or greater is automated. CAP scores have also been available since Botometer-V3. The CAP scores are Bayesian posteriors that reflect both the results from the classifier and prior knowledge of the prevalence of bots on Twitter, so as to balance false positives with false negatives. For example, suppose an account has a raw bot score of 0.96/1 (equivalent to 4.8/5 display score on the website) and a CAP score of 90%. This means that 90% of accounts with a raw bot score above 0.96 are labeled as bots, or, as indicated on the website, 10% of accounts with a bot score above 4.8/5 are labeled as humans. In other words, if you use a threshold of 0.96 on the raw bot score (or 4.8 on the display score) to classify accounts as human/bot, you would wrongly classify 10% of accounts as bots—a false positive rate of 10%. This helps researchers determine an appropriate threshold based on acceptable false positive and false negative error rates for a given analysis.

Fast bot classification

When Botometer-V4 was released, a new model called BotometerLite was added to the Botometer family [42]. BotometerLite was created to enable fast bot detection for large scale datasets. The speed of bot detection methods is bounded by the platform's rate limits. For example, the Twitter API endpoint used by Botometer-V4 to fetch an account's most recent 200 tweets and recent mentions from other users has a limit of 43,200 accounts per app key, per day. Many studies using Twitter data

have millions of accounts to analyze; with Botometer-V4, this may take weeks or even months.

To achieve scalability, BotometerLite relies only on features extracted from user metadata, contained in the so-called user object from the Twitter API. The rate limit for fetching user objects is over 200 times the rate limit that bounds Botometer-V4. Moreover, each tweet collected from Twitter has an embedded user object. This brings two extra advantages. First, once tweets are collected, no extra queries to Twitter are needed for bot detection. Second, the user object embedded in each tweet reflects the user profile at the moment when the tweet is collected. This makes bot detection on archived historical data possible.

In addition to the improved scalability, BotometerLite employs a novel data selection mechanism to ensure its accuracy and generalizability. Instead of throwing all training data into the classifier, a subset is selected by optimizing three evaluation metrics: cross-validation accuracy on the training data, generalization to hold-out datasets, and consistency with Botometer. This mechanism was inspired by the observation that some datasets are contradictory to each other. After evaluating the classifiers trained on all possible combinations of candidate training sets, the winning classifier only uses five out of eight datasets but performs well in terms of all evaluation metrics.

BotometerLite allows researchers to analyze large-volume streams of accounts in real time, while the limited training data may involve a compromise in accuracy on certain bot classes compared to Botometer-V4. In terms of how to choose between the two endpoints, we still recommend using Botometer-V4 when feasible since it analyzes more data and produces more detailed results.

Botometer interface

Although the machine learning model might seem complicated, the interface of Botometer is designed to be easy to use. Botometer has a website and API endpoints with similar functionality. The website⁵ is handy for users who need to quickly check several accounts. With a Twitter account, users can access the Botometer website from any web browsers, even on their mobile devices. The website is straightforward to use: after authorizing Botometer to fetch Twitter data, users just need to type a Twitter handle of interest and click the “Check user” button.

The Botometer Pro API⁶ can be more useful for research since it allows to programmatically check accounts in bulk. The API is hosted by RapidAPI, a platform that helps developers manage API rate limits and user subscriptions. Using the Botometer API requires keys associated with a Twitter app, which can be obtained through Twitter’s developer portal.⁷ One also needs a RapidAPI account and a subscription to one of the API usage plans.

⁵ botometer.org.

⁶ rapidapi.com/OSoMe/api/botometer-pro.

⁷ developer.twitter.com.

When querying the API, users are responsible to send the required data (i.e., 200 most recent tweets by the account being checked and tweets mentioning this account) in a specified format through HTTPS requests. The Botometer API will process the data and return the results. While queries can be sent through any programming language, we recommend using Python and the official `botometer-python` package that we maintain.⁸ The package can fetch data from Twitter, format the data, and query the API on behalf of the user with a few lines of code:

```
import botometer

bom = botometer.Botometer(
    rapidapi_key="XYZ",
    consumer_key="XYZ",
    consumer_secret="XYZ",
    access_token="XYZ",
    access_token_secret="XYZ"
)
result = bom.check_account("@yang3kc")

print(f"Bot score={result['display_scores']['english']['overall']/5}")
print(f"CAP score={result['cap']['english']:.2f}")
```

BotometerLite is also available as an endpoint through the Botometer Pro API. We list the the input, output, and limitations of the API endpoints for Botometer-V4 and BotometerLite side by side in Table 2. We also summarize the common resources for using Botometer in Table 3 to help the readers navigate these resources.

Note that both Botometer and Twitter APIs have rate limits, meaning that users can only make a certain number of queries in a given time period. Please check the respective websites for detailed documentation. Getting familiar with the rate limits can help researchers better estimate the time needed for their analysis.

Case study

Since some readers may not be familiar with programming, querying the API could be challenging. Moreover, analyzing the results returned by Botometer API is not trivial. In this section, we provide a simple case study as a demonstration. Different ways of analyzing the data are shown with recommended practice. We share the code for this case study in a public repository⁹ so that readers can use it as a template for their own research. Next we outline the data collection and analysis steps implemented in this software repository.

⁸ github.com/IUNetSci/botometer-python.

⁹ github.com/osome-ju/Botometer101.

Table 2 Comparison of Botometer-V4 and BotometerLite APIs

	Botometer-V4	BotometerLite
Endpoint	Check account	Check account in bulk
Query payload	User object, 200 most recent tweets, mentions	List of user objects and timestamps
Response	Raw bot scores, sub-scores, CAP scores, basic account information, etc.	BotometerLite scores
Daily number of accounts allowed*	43,200	~ 8.6 million
Corresponding botometer-python method(s)	check_account	check_accounts_from_tweets, check_accounts_from_user_ids, check_accounts_from_screen_names

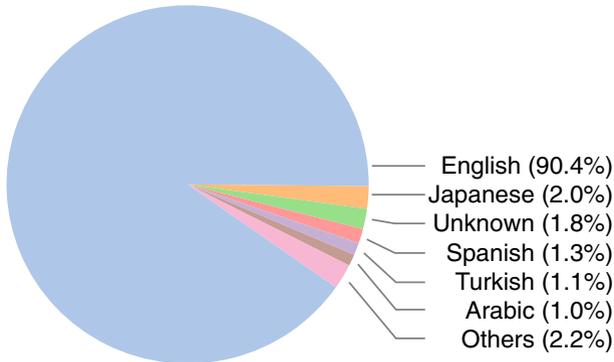
*The values represent the upper bounds based on Twitter's API rate limit when using a single app key. The actually numbers depend on other factors such as internet speed as well

Table 3 Common resources for using Botometer

Resource name	Resource	Note
Botometer website	botometer.org	Web interface of Botometer: useful for checking a small amount of accounts
Botometer Pro API	rapidapi.com/OSoMe/api/botometer-pro	API of Botometer: useful for checking accounts in bulk programmatically
Botometer-python package	github.com/IUNetSci/botometer-python	Python package to access Botometer Pro API
Botometer case study	github.com/osome-iu/Botometer101	Case study using Botometer with source code
Bot repository	botometer.osome.iu.edu/bot-repository	Annotated training datasets for Botometer

Table 4 Numbers of tweets and unique accounts mentioning different cashtags in raw data and analytical sample

Cashtag	Raw data		Analytical sample	
	Tweets	Unique accounts	Tweets	Unique accounts
\$SHIB	2000	1241	1819	1111
\$FLOKI	2000	937	1893	860
\$AAPL	2000	1107	1864	1006

**Fig. 2** Percentage of accounts using each language in the three datasets combined

Data collection

Let us consider two cryptocurrency cashtags, \$FLOKI and \$SHIB, and the cashtag of Apple Inc., \$AAPL, and attempt to quantify which is more amplified by bot-like accounts. A cashtag works like a hashtag but consists of a dollar sign “\$” and a stock or cryptocurrency symbol to help users track related discussions. We use Tweepy,¹⁰ a Python package that helps access the Twitter API, to search tweets containing these cashtags. For each cashtag, we only collect 2000 tweets, which are sufficient for the demonstration.

First, let us count the number of unique accounts in each dataset, as shown in Table 4. The number of unique accounts is much smaller than the number of tweets in all three datasets, suggesting that some accounts tweeted the same cashtag multiple times.

The next step is to query the Botometer API for bot analysis. Instead of going through each tweet and check every user encountered, researchers can keep a record of accounts already queried to avoid repetition and increase efficiency. The Botometer API returns rich information about each account. We recommend storing the full results from Botometer for flexibility.

As mentioned above, Botometer generates an overall score and a language-independent score. Since the two scores come from different classifiers, they are not

¹⁰ tweepy.org.

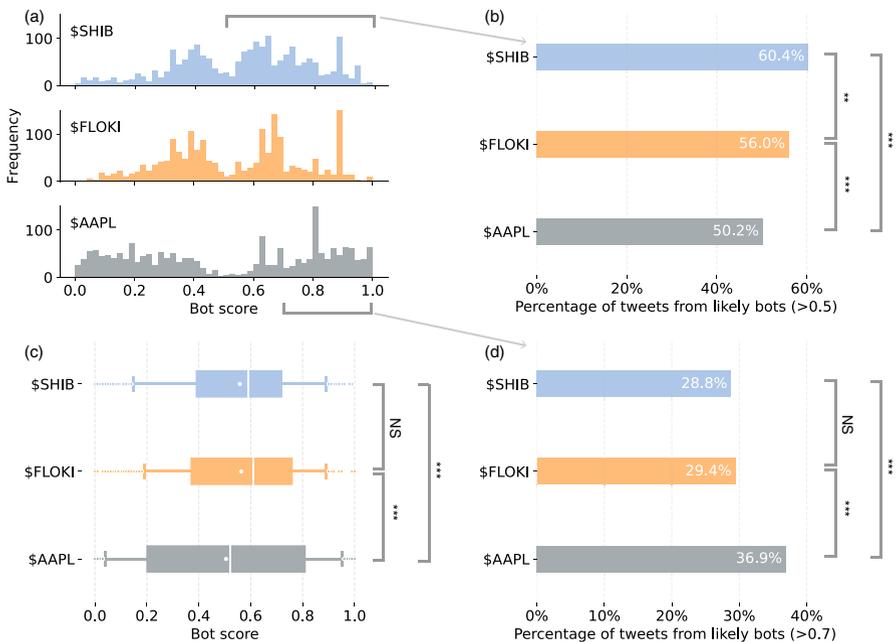


Fig. 3 **a** Bot score distributions for tweets mentioning different cashtags. **b** Percentage of tweets posted by likely bots using 0.5 as a threshold. **c** Box plots of the bot scores for tweets mentioning different cashtags. The white lines indicate the median values; the white dots indicate the mean values. **d** Similar to **b** but using a bot score threshold of 0.7. Statistical tests are performed for pairs of results in **b–d**. Significance level is represented by the stars: *** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$, NS= $p > 0.05$

comparable and should not be mixed together. To decide which one to use, let us calculate the proportion of accounts using each language. We can see in Fig. 2 that the majority of accounts in our raw data tweet in English. Therefore we only include English-speaking accounts and their tweets in our analytical sample (see Table 4 for summary statistics) and use the overall bot score.

Analysis

We plot the bot score distribution for tweets mentioning each cashtag in Fig. 3a. Here we base our analysis on the raw scores in the unit interval. Since we are interested in the bot activity level of each cashtag, we use tweets (as opposed to accounts) as the units of analysis. This means that accounts tweeting the same cashtag multiple times have a larger contribution.

In all three cases, the distribution has a bimodal pattern, a result of the ESC architecture of Botometer-V4. We can observe some spikes in all cases, which are caused by accounts tweeting the same cashtag repeatedly. For example, the spike near 0.89 for \$SHIB and \$FLOKI comes from a bot-like account that replied the same message promoting cryptocurrency tokens to a large number of tweets containing the keyword “NFT”; see the screenshot of the message in Fig. 4.

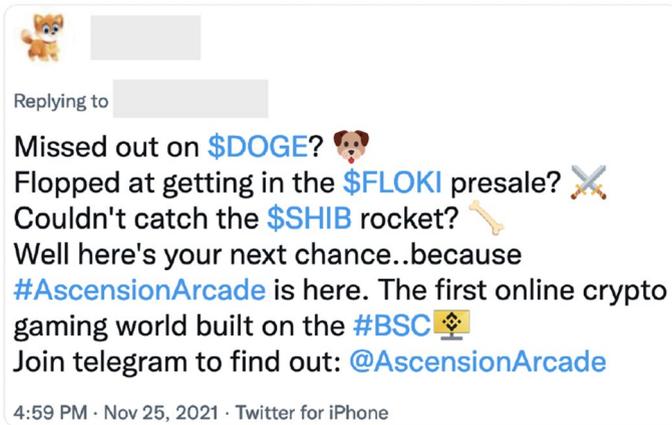


Fig. 4 Screenshot of a bot-like account replying to a tweet containing the keyword “NFT” with a message promoting cryptocurrencies. The same message was replied by this account to a large number of tweets

To address our research question, we need to quantify the bot activity level for each cashtag and compare them. The first approach is to compare their bot score distributions with two-sided Mann–Whitney U tests (see results in Fig. 3c). The bot score distributions of \$SHIB and \$FLOKI are not significantly different from each other ($p = 0.56$), but both of them have a higher bot activity level than \$AAPL (\$SHIB vs. \$AAPL: $p < 0.001$; \$FLOKI vs. \$AAPL: $p < 0.001$).

The second approach dichotomizes the bot scores and considers the accounts with scores higher than a threshold as likely bots. Then the proportion of tweets from likely bots can be calculated and compared. In this approach, a threshold has to be chosen. In the literature, 0.5 is the most common choice [2, 4, 37]; higher values, such as 0.7 [38] and 0.8 [13], are also used. One may also consider running the same analysis with different threshold values to test the robustness of the findings [2].

Here we use both 0.5 and 0.7 as thresholds and show the results in Fig. 3b, d, respectively. We apply two-proportions z -tests to estimate the significance level of the differences. When using 0.5 as the threshold, the percentage of tweets from likely bots that mentioned \$SHIB is significantly higher than those in the \$FLOKI ($p = 0.009$) and \$AAPL datasets ($p < 0.001$). The percentage of tweets from likely bots that mentioned \$FLOKI is also significantly higher than that in the \$AAPL dataset ($p < 0.001$). However, when using 0.7 as the threshold, the results change: percentages of tweets from likely bots in \$SHIB and \$FLOKI datasets are no longer significantly different from each other ($p = 0.38$); both of them are lower than that in the \$AAPL dataset (\$SHIB vs. \$AAPL: $p < 0.001$; \$FLOKI vs. \$AAPL: $p < 0.001$).

In other studies, different approaches or threshold choices may yield consistent results. However, they lead to seemingly different conclusions in this case. This is because different measures represent different properties of the bot score distribution. If we revisit Fig. 3a, we can see that although the distributions of \$SHIB and \$FLOKI scores have more mass in the (0.5, 1] region than that of \$AAPL scores, the

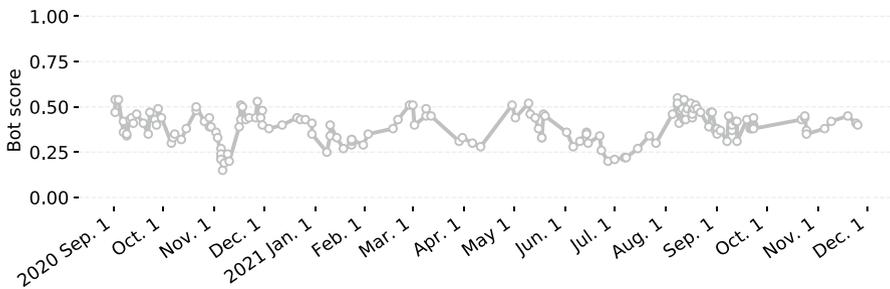


Fig. 5 Time series of bot scores of an account from September 2020 to November 2021. The queries were not made regularly, so the time intervals between consecutive data points vary

mass tends to concentrate around 0.6, while the distribution of \$AAPL scores has more mass near 1. This nuanced difference causes the contradictory results when using different threshold values.

By reconciling the results from different approaches, we can answer our research question now. It appears that discussions about the cryptocurrencies \$SHIB and \$FLOKI show more automated activities than that about \$AAPL, but among the accounts tweeting \$AAPL, we find more highly automated bot-like accounts. Note that the analysis here is mainly for demonstrating the use of Botometer; the samples of tweets analyzed are small and not representative of the entire discussion, so the conclusions only reflect the status of the collected data and should not be generalized.

Recommended practice

The sections above cover some recommended practice such as being careful when interpreting raw bot scores, being mindful about user language, and being aware of different versions of Botometer. Here we make a few more recommendations to help avoid common pitfalls.

Transient nature of Botometer scores

Recall that Botometer uses the 200 most recent tweets by an account and other tweets mentioning the account for analysis. This means that the results of Botometer change over time, especially for very active accounts. To demonstrate this, we plot the time series of the overall bot score of an account in Fig. 5. This account posts roughly 16 tweets each week and gets mentioned by others frequently. We can see that the bot score fluctuates over time. In some other cases, an account might be suspended or removed after a while, making it impossible to analyze.

Due to the transient nature of Botometer scores, a single bot score only reflects the status of the account at the moment when it is evaluated. Users should be careful when drawing conclusions based on the bot scores of individual accounts. For researchers, a common practice is to collect tweets first, then perform bot detection

later. To reduce the effect of unavailable accounts and to keep the bot scores relevant, bot analysis should be conducted right after data collection.

Evaluating bot score distributions

Whenever possible, we recommend collecting large datasets and use statistical analyses to evaluate bot activity based on comparisons of score distributions across different groups of accounts. As demonstrated in the case study, bot score distributions can reveal rich information about the data. Using distributions for analysis also reduces the uncertainty level of Botometer due to its imperfection and transient nature. Most importantly, comparing distributions of scores—e.g., for accounts tweeting about a given topic versus a suitable baseline—allows for statistical tests that are impossible at the level of individual accounts.

Validating thresholds

In some analyses, dichotomizing the bot scores based on a threshold is necessary. In these cases, we recommend validating the choice of threshold. For researchers with the ability and resources, the ideal approach is to manually annotate a batch of bot and human accounts in their datasets. Such a preliminary analysis could be used, first, to determine whether Botometer is a helpful tool to evaluate a given scenario. Assuming it is, one can then vary the threshold and select the value that optimizes some appropriate metric on the annotated accounts. Depending on the desire to maximize accuracy, minimize false positive errors, minimize false negative errors, or some combination, one can use metrics, such as accuracy, precision, recall, or F1. When annotating additional accounts is not feasible, we suggest running multiple analyses using different threshold choices to confirm the robustness of the findings.

Using Botometer in a civil way

We have noticed that Botometer has been used to attack others. For example, some users may call others with whom they disagree “bots” and use the results of Botometer as justification. This is a misuse of Botometer. Users should keep in mind that any classifier such as Botometer can mislabel individual accounts. Furthermore, even if an account is automated, it does not mean it is deceptive or malicious. Most importantly, such name calling is not helpful for creating healthy and informative conversations.

Acknowledgements We credit Onur Varol, Clayton A. Davis, Mohsen Sayyadiharikandeh, Pik-Mai Hui, and Alessandro Flammini for Botometer development and research. We are grateful to Chris Torres-Lugo, Prashant Shiralkar, Mohsen JafariAsbagh, Zoher Kachwala, Gregory Maus, James Caverlee, Kyumin Lee, Qiaozhu Mei, Zhe Zhao, Aram Galstyan, Shradha Gyaneshwar Baranwal, Andy Patel, and Josh Emerson for contributing insights and training data.

Funding The development of Botometer has been supported in part by DARPA (Grants W911NF-12-1-0037 and W911NF-17-C-0094), NSF (Grant CCF-1101743), ONR (Grant N15A-020-0053), AFOSR (Award FA9550-17-1-0327), NIH (Award 5R01DA039928-03), the James McDonnell

Foundation (Grant 220020274), Craig Newmark Philanthropies, and Knight Foundation. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Data Availability Statement Code used in this paper is available at: <https://github.com/osome-iu/Botometer101>. The datasets analysed in the present study are not publicly available due to Twitter's terms of service. However, readers can use the shared code to replicate the analysis by collecting data from Twitter on their own.

Declarations

Conflict of interest The authors declare that they have no competing interests.

References

1. Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, 59(7), 96–104.
2. Shao, C., Ciampaglia, G. L., Varol, O., Yang, K.-C., Flammini, A., & Menczer, F. (2018). The spread of low-credibility content by social bots. *Nature Communications*, 9(1), 4787.
3. Gorodnichenko, Y., Pham, T., & Talavera, O. (2021). Social media, sentiment and public opinions: Evidence from #Brexit and #USElection. *European Economic Review*, 136, 103772.
4. Bessi, A., Ferrara, E. (2016). Social bots distort the 2016 U.S. Presidential election online discussion. *First Monday*
5. Ferrara, E., Chang, H., Chen, E., Muric, G., Patel, J. (2020). Characterizing social media manipulation in the 2020 U.S. presidential election. *First Monday*
6. Ferrara, E. (2017). Disinformation and social bot operations in the run up to the 2017 French presidential election. *First Monday*
7. Bastos, M., & Mercea, D. (2018). The public accountability of social platforms: Lessons from a study on bots and trolls in the Brexit campaign. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2128), 20180003.
8. Bastos, M. T., & Mercea, D. (2019). The brexit botnet and user-generated hyperpartisan news. *Social Science Computer Review*, 37(1), 38–54.
9. Duh, A., Slak Rupnik, M., & Korošak, D. (2018). Collective behavior of social bots is encoded in their temporal twitter activity. *Big Data*, 6(2), 113–123.
10. Keller, T. R., & Klinger, U. (2019). Social bots in election campaigns: Theoretical, empirical, and methodological implications. *Political Communication*, 36(1), 171–189.
11. Stella, M., Ferrara, E., & Domenico, M. D. (2018). Bots increase exposure to negative and inflammatory content in online social systems. *Proceedings of the National Academy of Sciences*, 115(49), 12435–12440.
12. Jamison, A. M., Broniatowski, D. A., & Quinn, S. C. (2019). Malicious actors on twitter: A guide for public health researchers. *American Journal of Public Health*, 109(5), 688–692.
13. Broniatowski, D. A., Jamison, A. M., Qi, S., AlKulaib, L., Chen, T., Benton, A., et al. (2018). Weaponized health communication: Twitter bots and russian trolls amplify the vaccine debate. *American Journal of Public Health*, 108(10), 1378–1384.
14. Yuan, X., Schuchard, R.J., Crooks, A.T. (2019). Examining emergent communities and social bots within the polarized online vaccination debate in Twitter. *Social Media + Society* 5 (3), 2056305119865465
15. Ferrara, E. (2020). What types of COVID-19 conspiracies are populated by Twitter bots? *First Monday*
16. Shi, W., Liu, D., Yang, J., Zhang, J., Wen, S., & Su, J. (2020). Social bots' sentiment engagement in health emergencies: A topic-based analysis of the COVID-19 pandemic discussions on twitter. *International Journal of Environmental Research and Public Health*, 17(22), 8701.
17. Uyheng, J., & Carley, K. M. (2020). Bots and online hate during the COVID-19 pandemic: Case studies in the United States and the Philippines. *Journal of Computational Social Science*, 3(2), 445–468.

18. Yang, K.-C., Torres-Lugo, C., Menczer, F. (2020). Prevalence of low-credibility information on twitter during the COVID-19 outbreak. In *Proceedings of the ICWSM International Workshop on Cyber Social Threats*
19. Allem, J.-P., Escobedo, P., & Dharmapuri, L. (2020). Cannabis surveillance with twitter data: Emerging topics and social bots. *American Journal of Public Health, 110*(3), 357–362.
20. Marlow, T., Miller, S., Roberts, J.T. (2020). Twitter discourses on climate change: Exploring topics and the presence of bots. *SocArXiv*. <https://doi.org/10.31235/osf.io/h6ktm>
21. Marlow, T., Miller, S., Roberts, J.T. (2021). Bots and online climate discourses: Twitter discourse on President Trump’s announcement of U.S. withdrawal from the Paris Agreement. *Climate Policy, 21*(6), 765–777
22. Chen, C.-F., Shi, W., Yang, J., Fu, H.-H. (2021). Social bots’ role in climate change discussion on Twitter: Measuring standpoints, topics, and interaction strategies. *Advances in Climate Change Research*
23. Nizzoli, L., Tardelli, S., Avvenuti, M., Cresci, S., Tesconi, M., & Ferrara, E. (2020). Charting the landscape of online cryptocurrency manipulation. *IEEE Access, 8*, 113230–113245.
24. Cresci, S., Lillo, F., Regoli, D., Tardelli, S., & Tesconi, M. (2019). Cashtag piggybacking: Uncovering spam and bot activity in stock microblogs on twitter. *ACM Transactions on the Web, 13*(2), 11–11127.
25. Fan, R., Talavera, O., & Tran, V. (2020). Social media bots and stock markets. *European Financial Management, 26*(3), 753–777.
26. Keller, F. B., Schoch, D., Stier, S., & Yang, J. (2020). Political astroturfing on twitter: How to coordinate a disinformation campaign. *Political Communication, 37*(2), 256–280.
27. Bilton, N. (2014). Social media bots offer phony friends and real profit. *The New York Times*. <https://www.nytimes.com/2014/11/20/fashion/social-media-bots-offer-phony-friends-and-real-profit.html>
28. Confessore, N., Dance, G.J.X., Harris, R., Hansen, M.(2018). The follower factory. *The New York Times*. <https://www.nytimes.com/interactive/2018/01/27/technology/social-media-bots.html>, <https://www.nytimes.com/interactive/2018/01/27/technology/social-media-bots.html>
29. Varol, O., & Uluturk, I. (2020). Journalists on twitter: Self-branding, audiences, and involvement of bots. *Journal of Computational Social Science, 3*(1), 83–101.
30. Albadi, N., Kurdi, M., Mishra, S.(2019). Hateful people or hateful bots? Detection and characterization of bots spreading religious hatred in arabic social media. In *Proceedings of the ACM on Human-Computer Interaction 3*(CSCW), 61–16125
31. Caldarelli, G., De Nicola, R., Del Vigna, F., Petrocchi, M., & Saracco, F. (2020). The role of bot squads in the political propaganda on Twitter. *Communications Physics, 3*(1), 1–15.
32. Torres-Lugo, C., Yang, K.-C., Menczer, F. (2022). The manufacture of political echo chambers by follow train abuse on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*.
33. Chen, W., Pacheco, D., Yang, K.-C., & Menczer, F. (2021). Neutral bots probe political bias on social media. *Nature Communications, 12*, 5580.
34. Yang, K.-C., Varol, O., Davis, C. A., Ferrara, E., Flammini, A., & Menczer, F. (2019). Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Technologies, 1*(1), 48–61.
35. Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J.W., Christakis, N.A., Couzin, I.D., Jackson, M.O., Jennings, N.R., Kamar, E., Kloumann, I.M., Laroche, H., Lazer, D., McElreath, R., Mislove, A., Parkes, D.C., Pentland, A.S., Roberts, M.E., Shariff, A., Tenenbaum, J.B., Wellman, M. (2019). Machine behaviour. *Nature 568*(7753)
36. Ledford, H. (2020). Social scientists battle bots to glean insights from online chatter. *Nature, 578*(7793), 17–17.
37. Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science, 359*(6380), 1146–1151.
38. Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., Lazer, D. (2019). Fake news on Twitter during the 2016 U.S. presidential election. *Science, 363*(6425), 374–378
39. Bovet, A., & Makse, H. A. (2019). Influence of fake news in Twitter during the 2016 US presidential election. *Nature Communications, 10*(1), 7.
40. Varol, O., Ferrara, E., Davis, C.A., Menczer, F., Flammini, A. (2017). Online human-bot interactions: Detection, estimation, and characterization. In *Proceedings of the International AAAI Conference on Web and Social Media*

41. Sayyadiharikandeh, M., Varol, O., Yang, K.-C., Flammini, A., Menczer, F. (2020). Detection of novel social bots by ensembles of specialized classifiers. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*, pp. 2725–2732
42. Yang, K.-C., Varol, O., Hui, P.-M., & Menczer, F. (2020). Scalable and generalizable social bot detection through data selection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(01), 1096–1103.
43. Chavoshi, N., Hamooni, H., Mueen, A. (2016). Debot: Twitter bot detection via warped correlation. In: *ICDM*, pp. 817–822
44. Echeverria, J., Zhou, S. (2017). Discovery, retrieval, and analysis of the ‘star wars’ botnet in twitter. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 1–8
45. Cresci, S., Di Pietro, R., Petrocchi, M., Spognardi, A., Tesconi, M. (2017). 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*, pp. 963–972
46. Gilani, Z., Farahbakhsh, R., Tyson, G., Wang, L., Crowcroft, J. (2017). Of bots and humans (on Twitter). In *Proceedings of the International Conference on Advances in Social Networks Analysis and Mining*, pp. 349–354 . ACM
47. Mazza, M., Cresci, S., Avvenuti, M., Quattrociocchi, W., Tesconi, M. (2019). RTbust: Exploiting temporal patterns for botnet detection on twitter. In *Proceedings of the 10th ACM Conference on Web Science*, pp. 183–192
48. Cresci, S., Lillo, F., Regoli, D., Tardelli, S., Tesconi, M. (2018). \$FAKE: Evidence of spam and bot activity in stock microblogs on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, vol 12, p. 1
49. Rauchfleisch, A., Kaiser, J. (2020). Dataset for paper: The false positive problem of automatic bot detection in social science research. *Harvard Dataverse*. <https://doi.org/10.7910/DVN/XVCKRS/P2ZKRU>
50. Lee, K., Eoff, B.D., Caverlee, J. (2011). Seven months with the devils: A long-term study of content polluters on twitter. In *Proc. AAAI Intl. Conf. on Web and Social Media (ICWSM)*
51. Echeverria, J., Zhou, S. (2017). Discovery of the twitter bursty botnet. arXiv preprint [arXiv:1709.06740](https://arxiv.org/abs/1709.06740)
52. Rauchfleisch, A., & Kaiser, J. (2020). The False positive problem of automatic bot detection in social science research. *PLOS One*, 15(10), 0241045.
53. Martini, F., Samula, P., Keller, T. R., & Klinger, U. (2021). Bot, or not? Comparing three methods for detecting social bots in five political discourses. *Big Data & Society*, 8(2), 20539517211033570.
54. Dimitriadis, I., Georgiou, K., & Vakali, A. (2021). Social botomics: A systematic ensemble ml approach for explainable and multi-class bot detection. *Applied Sciences*, 11(21), 9857.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.