

How many bots are you following? ^{*}[†]

Alessandro Balestrucci^{1,2}

¹ IMT School for Advanced Study Lucca, Lucca, Italy

² Gran Sasso Science Institute, L'Aquila, Italy

`alessandro.balestrucci@gssi.it`

Abstract

Social Media are evolving as a pervasive source of news able to reach a larger audience through their spreading power. The main drawback is given by the presence of malicious accounts, known as social bots, which are often used to diffuse misleading information. Social bots are automated accounts whose goal is to interact with humans and influence them. Starting from the definition of *credulous* (i.e., human accounts with a high percentage of bot friends among their *followees*), in this work we aim to single out a regression model to derive, with an acceptable margin of error, the percentage of bot-followees of a human-operated account. The advantage lies in knowing, as a preventive measure, which users may be the target of bots' activities, hence more exposed to the misleading/unreliable content. Our results showed that the best regression model achieves a Mean Absolute Error of 3.62% and a Root Mean Squared Error of 5.96%, thus encouraging further research in this direction.

keywords – Credulous Twitter Users, Humans-Bots Interactions, Disinformation Spreading, Social Networks Analysis, Supervised Learning

1 Introduction

The widespread use of mobile devices, combined with the ease and cheapness of being connected, are some of the most important factors increasing the usage of Online Social Media (OSM) as means of communication [35]. Several studies [11, 34, 3] recently confirmed OSM (such as Facebook and Twitter) as more pervasive than traditional mass media (e.g., radio, newspaper, etc.), especially among youngest people [35].

As consequence of the OSM's pervasiveness, a larger audience is reached, but several issues arise about the veracity of the circulating news. This leads to fake news and misinformation that are intentionally propagated for harmful purposes [27]. The information diffusion on OSM is often supported by automated accounts controlled, totally (bots) or partially (cyborg), by ad-hoc software applications [45]. Under fictive identity, such automated accounts actively interact (social bots [18]) with genuine users and share/produce contents of doubtful credibility. Governments¹, academics [50] and OSM administrators² are struggling to face these problems.

^{*}Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

[†]Partially supported by the European Unions Horizon 2020 programme (grant agreement No. 830892, SPARTA) and by IMT Scuola Alti Studi Lucca: Integrated Activity Project TOFFEE TOOLS for Fighting FakeEs. We would like to thank my supervisor Rocco De Nicola and Catia Trubiani for the useful discussions and suggestions regarding this work.

¹<https://tinyurl.com/yym2xa3v>

²<https://tinyurl.com/yac71sn6>

Despite the efforts spent by OSM administrators in removing suspicious accounts³⁴, and by researchers in improving bot detection techniques [33, 2], this plague is far from being eradicated. In fact, by recent estimates on Twitter, the percentage of bots among the active users ranges between 9% and 15% [45]. Furthermore, a new generation of bots, more sophisticated than the previous ones, are able to avoid detection by mimicking human behaviour [14].

In some cases, bots have been used to call volunteers in case of emergencies [40] or to spread academic events such as conferences, but these are only exceptions. In fact, a dominant use of these artificial entities is for malicious purposes, e.g., to encourage hate speeches, misconception and, more in general to influence people [44, 25].

The effectiveness of such malicious activities, intended to manipulate public opinion, has been investigated on the Brexit referendum [22], the US Presidential election 2016 [8], the elections in France [17] and Mexico [9]. These are few well-known examples of how and to what extent bots and disinformation can damage democracy, especially if this kind of campaigns target specific categories of users.

A recent study stated that the majority of genuine (human) users usually do not check the reliability of contents on OSM, and many even share these articles [16]; hence contributing, although unknowingly, to diffuse unreliable content. Moreover, in [26] models for influence propagation in OSM have been studied, and a strong correlation emerged between the target nodes (to influence) and the role of their neighbours (social contacts).

Using Twitter as benchmark, this paper focuses on finding a way to determine, as precisely as possible, the amount of potential malicious accounts (bots) a human is following (*bot-followees*). In our opinion, the more bots a user is following, the more she/he is exposed to potentially malicious activities. Specifically, by using Machine Learning (ML) techniques, we aim to single out a regression model able to predict, for a human-operated account, the percentage of its *bot-followees*. Starting from a publicly available dataset⁵ of *credulous* users, built in [7], and by using commonly employed features from the literature [45, 12], we build several regression models by means of various ML algorithms. Since this work deal with a regression task, we cannot use the well-known metrics usually used for classification tasks (e.g., accuracy, F-score); instead, as in [32], we will use the Mean Absolute Error (MAE) and the Root Mean Squared Error (RMSE) for evaluation.

For the sake of completeness, we experiment on two “versions” of the considered dataset: firstly by considering the subset of *credulous* users only (with 316 instances, henceforth called *credulous-only*), and then on the whole set of human-operated accounts (with 2,838 instances, named *all humans*). The experimental results are promising. The best model trained by using *credulous-only*, achieves a MAE of 4.32% and a RMSE of 6.62%. On the other hand, astonishingly, the best model trained on the whole humans set, fits better; in fact it obtains a MAE of 3.62% and a RMSE of 5.96%. Although RMSE scores are not the best obtained, by following the finding in [48], we prefer to give more importance to MAE scores than to RMSE.

Despite these encouraging results, additional efforts are needed to improve users’ awareness. The usefulness in improving the accuracy of these models is twofold. Firstly, in a preventive way, to identify sensible users more exposed to the malicious activities of bots and by being identified as good targets of misinformation campaigns. Secondly, to safeguard the usefulness, credibility and effectiveness of social media as mean of communication. Moreover, by adapting the concept of users’ *following* to other types of relationships between users, it is possible to conduct these same experiments in other Social Media, e.g. Facebook.

³Facebook: <https://tinyurl.com/y3yzvpah>

⁴Twitter: <https://tinyurl.com/y3efs8s5>

⁵Dataset: <https://tinyurl.com/y4o98c71>

The remainder of this paper is organized as follow: Section 2 presents the related work highlighting differences and novelty of our approach; Section 3 exposes the general approach and provides the details of the performed experiments; Section 4 illustrates the experimental results; Section 5 discusses the main findings, and Section 6 concludes the paper by outlining future research.

2 Related Work

This paper aims to investigate the dynamics and causes of human users influenced by Social Media, especially if interacting with bots and exposed to misinformation. In the following, we review the related work dealing with similar goals.

Despite this cannot be considered an exhaustive literature review, we tried to supply concepts, definitions and adopted strategies related to interactions between human-operated accounts and bots on Social Media; and to stress the differences and novelty in this work.

In [46] the authors stress the importance of adopting protection mechanisms in Social Media to protect human-operated accounts. To this purpose, the authors provided the definition of *susceptible* users, i.e., users who start to interact with a social bot (even if only once). By introducing a set of 97 features (categorized in *linguistic*, *network* and *behavioural*), a binary classifier is successfully produced to spot out *susceptible* users. A regression model was also built to predict users level of *susceptibility* without success. Inspired by the definition of *susceptible* users, in our work we investigate on the challenging task to build a regression model able to predict the percentage of *bot-followees* of human-operated accounts.

On the same line of the described research, in [47] the authors conducted a feature study to single out those features useful to predict whether a user is likely to interact with a bot. To this end, the considered users have been contacted by a bot through a tweet mention, and when replying users are labelled as susceptible. The drawback of this approach is represented by the key feature (i.e., *klout* score), no longer available due to the related web-service closure.

In [31] a comprehensive categorization scheme for social bot attacks in Twitter has been proposed. The analysis of each attack is considered along different dimensions, i.e., *targets*, *account types*, *vulnerabilities*, *attack methods* and *results*. This categorization is useful to investigate to which class of attacks users with a considerable percentage of *bot-followees* are vulnerable. Specifically, the authors observed the impact of social bots in link creation between targeted human-operated accounts in Twitter.

Inspired by [46], in [42] *gullible* users have been defined in relation to their susceptibility to fake news. The authors present five degrees of susceptibility (referred to a user’s reply to a fake news) and, by using the same features of [46], aim to build a multi-class classifier predicting the susceptibility level. The classifier achieves an AUC of 0.82. Similarly to this approach, we also used existing feature sets (Botometer [15] and ClassA [12]), but applied to a regression task.

An interesting case study is presented in [39], where human’s skills have been tested in recognizing fake accounts generated by the M3 app [29]. In general, there is a good humans’ ability to distinguish the generated accounts, but there are no details on the characteristics of such fake accounts. Sharing the same motivation, our work aims to produce a regression model that, complementary to [39], quantifies when a human user is not able to recognize a bot.

We conclude this section by mentioning our previous works [6, 7]. Beside the definitions of *susceptible* users [46] and *gullible* users [42], in [6] we introduced the concept of *credulous* users, i.e., human operated accounts following a considerable amount of bots (i.e., *bot-followees* using Twitter’s slang). A set of rules have been introduced to discern whether a user is a credulous, and these rules allow to rank human operated accounts by relying on the ratio of bots (supported

by a bot detector) over their followees. On a dataset with more 700 human operated accounts, only 64 have been identified as credulous⁶. Here, the drawback is represented by the huge amount of data needed to investigate the followees of humans under analysis.

To overcome this issue, in [7], we built a binary classifier to find out credulous Twitter users by considering a larger *ground-truth* (of almost 3k⁷) and considering only features of the profile and not of users' followees. We ended up with a lightweight (in terms of costs for gathering the data) classifier. The achieved classification performance have been very promising with an accuracy of 93.27% and an *AUC* (Area Under the ROC curve) of 0.93.

Although this work is in line with the themes of the previous ones [6, 7], it differs in scope and strategy. The primary target is the prediction of a percentage value of bot-followees, and the adopted strategy is no longer the classification, but the regression.

3 Experimental Setup

This section explains the dataset, the features and the experimental design; specifying the metrics used to evaluate the performance of the trained models.

3.1 Dataset

The considered dataset includes 2,838 IDs of human-operated accounts on Twitter, presented in [7] and publicly available⁷. A label (0 or 1) is associated to each entry to indicate whether an account is *credulous* (with 1) or not (with 0).

For each account we downloaded all the information (timeline, tweet mentions and profile data) of its *followees*. Then, through of a bot detector [7], it has been possible to assign, to each *followee*, a label indicating whether it is a bot. This allows us to derive a ground-truth where each accounts is assigned to a percentage of its bot-followees (publicly available⁸). As anticipated in Section 1, two versions have been built: one with 316 entries labeled as credulous (*credulous-only*) and another one including all 2,838 human-operated accounts (*all_hums*). All accounts in our dataset belong to human users and taken from three datasets publicly available⁹ described below:

cresci-2015 [12]. This repository consists of three smaller datasets: (i) 469 Twitter accounts certified as human-operated during in a research project named *@TheFakeProject*. (ii) 1,481 accounts selected as genuine users as the result of a sociological study relying on manual verification. (iii) 833 fake accounts, bought from three different Twitter accounts online markets.

cresci-2017 [13]. This repository contains 3,474 Twitter accounts certified as humans and 6,609 social spambots (e.g., spammers of job offers and advertising products on sale at Amazon). The humans were selected through a hybrid crowd-sensing approach [5]; the authors of [13] randomly contacted Twitter users and asked simple questions in natural language, all replies were then manually verified.

varol-2017 [45]. This repository consists of 2,573 Twitter accounts selected after a manual annotation based on inspecting the profile details and the produced content. Overall, 1,747 Twitter accounts were annotated as human-operated and 826 as bots.

⁶First Credulous Dataset: <https://tinyurl.com/y61od2yz>

⁷Dataset: <https://tinyurl.com/y4o98c71>

⁸Used *ground-truth*: <https://tinyurl.com/tcjmbu>

⁹BotRepository:<https://tinyurl.com/yxfmqac>

3.2 Features

The features employed to represent the human accounts can be grouped in three sets. The first set inherits the features from the Botometer web service¹⁰, plus some additional ones introduced in [6]. We use the Botometer *macro categories scores*, which encompasses the essence of all the features (around 1,000) [15], on which the outputs’ scores are computed. We extend this set by considering also: the CAP (Completely Automated Profile) value, the “Scores” [49], the number of tweets and mentions. This augmented set of features will be called *Botometer+*.

The second set inherits the so-called *ClassA* features, singled out in [12]. We discard the one considering *duplicated pictures*, because we could not verify whether the picture of a profile was also used in another one. This subset is called *ClassA-*. We want to stress that these features can be calculated by only looking at the social profile of an account, without considering the timeline, that is instead considered in *Botometer+*. A schematic representation of these two feature sets is reported in table 5 in Appendix A.1.

The third set of features has been built through the union of the two previous ones and called *ALL_features*. We want to stress the dataset’s independency of the used feature sets that have been successfully used for binary classification task in [7]. The introduction of new features is out of the scope of this work.

3.3 Experimental Design

The experimental session starts by setting up the data; precisely, by transforming the entries of our dataset accordingly to the three sets of features.

Afterwards, to train regression models, 14 algorithms have been employed. The machine learning framework utilized to conduct the experiments is Weka [21]. It is worth to notice that, by default, the results obtained by the *experimenter* in the weka framework are cross-validated. Accordingly to the notation adopted by the tool, the algorithms are: **ZeroR**¹¹ (used to obtain a baseline value against which to compare the values of the other model), REPTree [37], Linear-Regression [28], k-Nearest Neighbour (IBk) [1], LWL [4], AdditiveRegression [20], Regression-ByDiscretization [19], M5Rules [38], DecisionStump [24], GaussianProcess [30], SMOreg [43], MultilayerPerceptron [36], MLPRegressor¹², RandomForest [10]. To investigate whether regression models, built on *credulous* users work better than the general case, we experiment on both versions (*credulous-only* and *all_hums*).

Results evaluation At best of our knowledge, there are few papers addressing this problem [46, 42]. Hence, no well-defined baseline is available from the literature to compare our results. To overcome this issue, we compare the performance of the trained models (obtained from the aforementioned algorithms) with the score calculated by **ZeroR** method, since it predicts, for each human-operated account, the average value of the population.

To evaluate models’ performance, two metrics have been used: Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) that are widely used for regression tasks¹³ [32]. The former measures the average of errors in a set of numerical predictions, between the real values and the predicted ones ($MAE = \frac{1}{n} \sum_{j=1}^n |y_{real_j} - y_{pred_j}|$); the latter measures the error too, but it stresses more the prediction error by raising the square of the difference between the real values and the predicted ones ($RMSE = \sqrt{\frac{1}{n} \sum_{j=1}^n (y_{real_j} - y_{pred_j})^2}$).

¹⁰Botometer web service (RapidAPI): <https://tinyurl.com/yyt282s>

¹¹ZeroR weka: <https://tinyurl.com/y4hdhp54>

¹²MLPRegressor weka: <https://tinyurl.com/y5krc6d2>

¹³<https://tinyurl.com/yd9ljcmj>

Through the experimental results analyzer embedded in Weka [41], we performed statistical tests (paired T-Test [23] with $\alpha = 0.05$) to determine which algorithm performs significantly better than the baseline, relatively to each set of features.

All the results are presented in the following section 4 and then discussed in section 5.

4 Experimental Results

For sake of clarity, Section 4.1 reports the results referring to the experiments performed on the set of credulous users only (*credulous-only*). Section 4.2 shows the outcome derived by considering all the human-operated accounts (*all_hums*).

Tables 1-4 have the same structure, differing in evaluation metric and ground-truth data entries. In the first column the algorithms mentioned in Section 3.3 are listed. The remaining columns show the scores obtained when dataset’s instances are represented according to a specific feature set (see section 3.2) namely *Botometer+*, *ClassA-* and *All_features*. The first row of the table contains the baseline obtained through the *ZeroR* method. The star symbol indicates some tables’ entries whose associated value is significantly lower than the baseline, according to the paired t-test performed by Weka (see results evaluation in Section 3.3). The lowest score is reported in bold.

4.1 Credulous-only

Table 1 reports the scores related to the RMSE. The baseline is offered by *ZeroR* showing a score equal to 6.73%. Note that in the column headed *Botometer+*, there are no starred values. Despite some of them are lower than the baseline (i.e., *LinearRegression*, *M5Rules*, *GaussianProcesses*, *SMOreg* and *RandomForest*), their values have not been considered significantly lower by Weka, in fact they are slightly better than the baseline. Among all values, the lowest one belongs to *GaussianProcesses* (6.48%) followed by *LinearRegression* (6.52%).

Algorithms	Feature sets		
	<i>Botometer+</i>	<i>ClassA-</i>	<i>All_features</i>
→ ZeroR ← (baseline)	6.73	6.73	6.73
REPTree	6.92	6.86	6.86
LinearRegression	6.52	8.52	8.62
IBk	8.71	8.95	8.02
LWL	6.84	6.10*	6.26
AdditiveRegression	6.79	6.30	6.20
RegressionByDiscretization	8.43	7.47	7.78
M5Rules	6.53	9.20	7.44
DecisionStump	6.90	6.15*	6.15*
GaussianProcesses	6.48	7.34	7.55
SMOreg	6.62	7.70	7.71
MultilayerPerceptron	10.79	11.97	11.82
MLPRegressor	7.59	7.50	6.86
RandomForest	6.60	6.15*	6.21

Table 1: Root Mean Squared Error (RMSE) on credulous users

Differently from the *Botometer+* case, the column *ClassA-* there are some starred values.

Algorithms	Feature sets		
	<i>Botometer+</i>	<i>ClassA-</i>	<i>All_features</i>
→ ZeroR ← (baseline)	4.84	4.84	4.84
REPTree	4.91	4.78	4.77
LinearRegression	4.68	5.18	5.19
IBk	5.88	6.17	5.44
LWL	4.83	4.36*	4.40*
AdditiveRegression	4.83	4.55	4.39*
RegressionByDiscretization	5.88	5.22	5.54
M5Rules	4.68	5.18	4.90
DecisionStump	4.83	4.36*	4.36*
GaussianProcesses	4.66	4.90	4.95
SMOreg	4.32*	4.64	4.67
MultilayerPerceptron	6.91	8.17	7.92
MLPRegressor	5.16	5.19	4.78
RandomForest	4.86	4.54	4.44

Table 2: Mean Absolute Error (MAE) on credulous users

Both *DecisionStump* and *RandomForest* achieve a RMSE score of 6.15%; but, the lowest value of 6.10% belongs to *LWL*, which also is the best score reported in Table 1.

The last column of Table 1 contains only one starred value which also the column’s lowest score (*DecisionStump*). Like for the previous column, also here there are some values lower than the baseline (*AdditiveRegression* with 6.20% and *RandomForest* with 6.21%), but not enough to be labeled as star entries.

Table 2 exposes the scores related to MAE. The indicated baseline is of 4.84%.

Looking at *Botometer+* features, the values lower than the baseline are: 4.83% (*LWL*, *AdditiveRegression*, *DecisionStump*), 4.68% (*LinearRegression*, *M5Rules*), 4.66% (*GaussianProcesses*) and 4.32% (*SMOreg*, starred). The latter is the lowest MAE.

When considering *ClassA-* features, the values lower than the baseline are: 4.78% (*LinearRegression*), 4.64% (*SMOreg*), 4.55% (*AdditiveRegression*), 4.54% (*RandomForest*) and 4.36% (*DecisionStump* and *LWL*, both starred).

If all features are analyzed, the values lower than the baseline are: 4.78% (*MLPRegressor*), 4.77% (*REPTree*), 4.67% (*SMOreg*), 4.44% (*RandomForest*), 4.40% (*LWL*, starred), 4.39% (*AdditiveRegression*, starred), 4.36% (*DecisionStump*, starred).

4.2 All humans

In this Section we report the results related to the experiments performed on all human-operated accounts (2,838 instances).

Table 3 reports the values concerning RMSE, with a baseline value of 6.25%.

In the second column (*Botometer+*), almost all the RMSE scores are lower than the baseline and starred, with the exception of *IBk* (7.73%), *RegressionByDiscretization* (6.32%) and *MultilayerPerceptron* (7.67%). With a score of 5.77%, *LinearRegression* has the lowest column’s RMSE (and the second better one in Table 3). Concerning the third column (*ClassA-*), the situation is slightly worse. Despite nine values have better scores than the baseline, only three of them are significantly lower: 6.02% (*REPTree*, the lowest) and 6.06% (*DecisionStump* and

Algorithms	Feature sets		
	<i>Botometer+</i>	<i>ClassA-</i>	<i>All_features</i>
→ ZeroR ← (baseline)	6.25	6.25	6.25
REPTree	5.96*	6.02*	5.93*
LinearRegression	5.77*	6.14	5.80*
IBk	7.73	8.58	7.59
LWL	5.91*	6.08	5.99*
AdditiveRegression	5.84*	6.07	5.80*
RegressionByDiscretization	6.32	6.43	6.83
M5Rules	6.02*	6.16	5.84*
DecisionStump	5.96*	6.06*	6.02*
GaussianProcesses	5.79*	6.13	5.83*
SMOreg	5.91*	6.36	5.96*
MultilayerPerceptron	7.67	6.53	9.47
MLPRegressor	5.89*	6.06*	6.84
RandomForest	5.92*	6.09	5.72*

Table 3: Root Mean Squared Error (RMSE) on all the human-operated accounts

MLPRegressor both). The forth column (*ClassA-*) shows a situation close to *Botometer+*'s case. In fact, with exception of four cases (i.e., *IBk*, *RegressionByDiscretization*, *MultilayerPerceptron* and *MLPRegressor*), all the other entries have significantly better values w.r.t. baseline (starred table's entries). The lowest RMSE of 5.72% is achieved by using *RandomForest* and it is the better score of Table 3.

Finally, Table 4 presents the MAE outcome. The calculated baseline is 4.21%. At first

Algorithms	Feature sets		
	<i>Botometer+</i>	<i>ClassA-</i>	<i>All_features</i>
→ ZeroR ← (baseline)	4.21	4.21	4.21
REPTree	3.95*	3.94*	3.87*
LinearRegression	3.84*	4.08	3.83*
IBk	5.07	5.43	4.95
LWL	3.97*	4.00*	3.98*
AdditiveRegression	3.89*	3.93*	3.76*
RegressionByDiscretization	4.16	4.24	4.36
M5Rules	3.91*	3.96*	3.82*
DecisionStump	4.06	3.99*	4.07
GaussianProcesses	3.87*	4.09	3.87*
SMOreg	3.67*	3.84*	3.62*
MultilayerPerceptron	4.90	4.39	5.14
MLPRegressor	3.88*	3.93*	4.07
RandomForest	3.96*	3.96*	3.77*

Table 4: Mean Absolute Error (MAE) on all the human-operated accounts

sight, regardless to the feature sets, almost all the entries are lower than baseline; and most of them are starred. When considering *Botometer+* features, the exceptions are: *IBk* (5.07%) and

MultilayerPerceptron (4.90%); *RegressionByDiscretization* (4.16%) and *DecisionStump* (4.06%) are not lower enough to gain the star. For *Botometer+*, the lowest MAE is 3.67% (the second better value) achieved by the model built by mean of *SMOreg*.

Similarly, when analyzing *ClassA-* features, the values higher than the baseline are: 5.43% (*IBk*), 4.24% (*RegressionByDiscretization*) and 4.39% (*MultilayerPerceptron*). The remaining values are significantly lower than the one in the *ZeroR*'s row but 4.08% (*LinearRegression*) and 4.09% (*GaussianProcesses*). Even in this case, the lowest score (3.84%) belongs to *SMOreg*. Like in the previous cases, even when all features are taken into account, almost all values are lower, not only compared to the baseline, but also compared (by row) to the values of the other two feature sets. The values overcoming the baseline are: 5.14% (*MultilayerPerceptron*), 4.95% (*IBk*) and 4.36% (*RegressionByDiscretization*). Except *MLPRegressor* and *DecisionStump* (both 4.07%), all the other entries have significantly better MAE values (starred). Once again, *SMOreg* outperforms other algorithms' scores for all features, with a MAE of 3.62%; the lowest in Table 4.

5 Discussion

It is possible to immediately notice that better results can be obtained when the full version of the dataset (*all_hums*) is considered rather than *credulous-only*, in terms of both quantity (number of models with lower performance than the baseline) and quality (statistical significance of the values).

However, focusing on baseline values of each metric, the highest ones are observed in the version *credulous-only*. This is due to a higher distance between the real values (by *credulous* class attribute) and the average value calculated on them. The fact that *all_hums*'s baselines have lower values indicates a better closeness to the average value of not-credulous instances.

Considering the MAE metric, the model generated by the *SMOreg* algorithm is the most accurate, regardless of the dataset versions. This concordance is limited exclusively to the algorithm, as these performances have been obtained with two different feature sets: 4.32% by using *Botometer+* for *credulous-only* version and 3.62% with *All_features* for the *All_hums* version.

But, considering the dataset with all instances, and looking the MAE's score obtained by *SMOreg* (using *Botometer+*'s features), we can see that the value is very similar (3,67%, the second-best result). This 0,05% loss (3,62 vs 3,67) can still be overshadowed by the advantage of not having to calculate the *ClassA-* features (included in *ALL_features*). Therefore, at least as far as MAE metrics are concerned, a representation in *Botometer+* features combined with the use of the *SMOreg* algorithm can be considered the best choice.

Good performances are also achieved by the algorithms *LWL* and *RandomForest*, producing the models with the best RMSE in *credulous-only* and *all_hums*, respectively. *DecisionStump* and *AdditiveRegression* produce good models too, and many times even with errors significantly lower than the baseline. As for the models generated through *REPTree*, *LinearRegression*, *M5Rules*, *GaussianProcesses* and *MLPRegressor*, it is possible to notice an inequality depending on the version of the dataset to which they are applied, proving more effective when used on the full version. Unfortunately, the use of *IBk*, *RegressionByDiscretization* and *MultilayerPerceptron* has not been profitable, regardless of both the set of used features and the version of the considered dataset.

In some cases, on the same algorithm, the score obtained by using *All_features* is identical (or very similar) to the one in *ClassA-* or *Botometer+*. These situations occur when the algorithm

prefers to consider the components of a certain feature set. Some examples are given by: *DecisionStump* in Tables 2 and 1, *GaussianProcesses* in Table 4 and *REPTree* in Table 1.

Further findings can be provided by studying, for each evaluation metric, to what extent the choice of a feature set (w.r.t. another one) can affect the overall performance. Regardless of the cost to calculate a feature set, the experimental results do not show a great disparity in preferring one feature set over another. Therefore, as far as the "MAE" is concerned, the previous assertion of preferring the *Botometer+* functionalities remains valid.

On the contrary, considering RMSE, the situation is more complicated. The best values come from the use of two distinct feature sets: *ClassA-* (*credulous-only*) and *All_features* (*all_hums*). By using a similar logic adopted for the MAE, it is possible to notice that, the second-best RMSE value for *credulous-only* is identical both for *ClassA-* and *All_features*. Therefore, unlike MAE case, if we concern about the RMSE only, we can prefer *All_features*, making the calculus of *ClassA-*'s features mandatory.

6 Conclusion

The recent literature shows a great deal of concern regarding the effectiveness of malicious bots and in their role to affect several domains (economics, politics, etc.). Not only as regards the spreading of fake news, but especially on their ability to interact with and deceive human users. In fact, the dissemination of low-credibility content is not only performed by social bots but also emphasized by genuine users by sharing it [16].

In this work, we devote our attention to single out a method that enables us to quantify, as precisely as possible, the percentage of bots followed (*bot-followees*) by human operated accounts in Twitter. By mean of machine learning regression techniques, we derived predictive models. The experimentation has been conducted firstly by considering *credulous* users only (i.e., humans with high percentage of bots among their friends) and then on all the human-operated accounts in the dataset.

For our experimentation we considered three feature sets, namely: *Botometer+*, *ClassA-* and their union *All_features*. We recall that *Botometer+*'s features are obtained by using the related web-service, and for each user the following input is required: profile data, tweets, mention-tweets. The calculus of *ClassA-*'s features is less expensive because such features are derived by analyzing users' data profile only. The best regression model, obtained on the complete dataset trained with the *SMOreg* algorithm, achieves promising performances, in fact it shows a MAE of 3.62% by using the feature set *All_features* and 3.67% with *Botometer+*'s features.

We think that having an estimate of how many bots a human account follows, can be a first step to the fact-checking of what people reads on their dashboards and from whom they are reading. Additionally, such approaches can help researchers and social media administrators to (i) identify potential targets of misinformation campaigns in advance and (ii) contributing to increase the usefulness, credibility and effectiveness of social media.

Further efforts have to be devoted to this research topic; as future work it can be interesting to focus on the nature and quality of the information that humans with a large percentage of *bot-followees* create and/or contribute to diffuse.

References

- [1] David W Aha, Dennis Kibler, and Marc K Albert. Instance-based learning algorithms. *Machine learning*, 6(1):37–66, 1991.

- [2] Shehu Amina, Raul Vera, Tooska Dargahi, and Ali Dehghantanha. A bibliometric analysis of botnet detection techniques. In *Handbook of Big Data and IoT Security*, pages 345–365. Springer, 2019.
- [3] Zeenab Aneez, Taberez Ahmed Neyazi, Antonis Kalogeropoulos, and Rasmus Kleis Nielsen. Reuters institute india digital news report. *Reuters Institute for the Study of Journalism/India Digital News Report*. Retrieved from https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2019-03/India_DNR_FINAL.pdf on, 26:19, 2019.
- [4] C. Atkeson, A. Moore, and S. Schaal. Locally weighted learning. *AI Review*, 1996.
- [5] Marco Avvenuti, Salvatore Bellomo, Stefano Cresci, Mariantonietta Noemi La Polla, and Maurizio Tesconi. Hybrid crowdsensing: A novel paradigm to combine the strengths of opportunistic and participatory crowdsensing. In *Proceedings of the 26th International Conference on World Wide Web Companion, Perth, Australia, April 3-7, 2017*, pages 1413–1421, 2017.
- [6] Alessandro Balestrucci, Rocco De Nicola, Omar Inverso, and Catia Trubiani. Identification of credulous users on twitter. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*, pages 2096–2103. ACM, 2019.
- [7] Alessandro Balestrucci, Rocco De Nicola, Marinella Petrocchi, and Catia Trubiani. Do you really follow them? automatic detection of credulous twitter users. *arXiv preprint arXiv:1909.03851 - to appear in Proceedings of IDEAL 2019*, 2019.
- [8] Alessandro Bessi and Emilio Ferrara. Social bots distort the 2016 us presidential election online discussion. *First Monday*, 21(11-7), 2016.
- [9] Samantha Bradshaw and Philip Howard. Troops, trolls and troublemakers: A global inventory of organized social media manipulation. *Oxford Internet Institute*, 2017.
- [10] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [11] Irene Costera Meijer and Tim Groot Kormelink. Checking, sharing, clicking and linking: Changing patterns of news use between 2004 and 2014. *Digital Journalism*, 3(5):664–679, 2015.
- [12] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. Fame for sale: Efficient detection of fake Twitter followers. *Decision Support Systems*, 80:56–71, 2015.
- [13] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In *Proceedings of the 26th International Conference on World Wide Web Companion, Perth, Australia, April 3-7, 2017*, pages 963–972, 2017.
- [14] Florian Daniel, Cinzia Cappiello, and Boualem Benatallah. Bots acting like humans: understanding and preventing harm. *IEEE Internet Computing*, 23(2):40–49, 2019.
- [15] Clayton Allen Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. Botornot: A system to evaluate social bots. In *Proceedings of the 25th International Conference Companion on World Wide Web*, pages 273–274. International World Wide Web Conferences Steering Committee, 2016.
- [16] Phillip George Efthimion, Scott Payne, and Nicholas Proferes. Supervised machine learning bot detection techniques to identify social twitter bots. *SMU Data Science Review*, 1(2):5, 2018.
- [17] Emilio Ferrara. Disinformation and social bot operations in the run up to the 2017 french presidential election. *First Monday*, 22(8), 2017.
- [18] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. The rise of social bots. *Communications of the ACM*, 59(7):96–104, 2016.
- [19] Eibe Frank and Remco R Bouckaert. Conditional density estimation with class probability estimators. In *Asian Conference on Machine Learning*, pages 65–81. Springer, 2009.
- [20] J.H. Friedman. Stochastic gradient boosting. Technical report, Stanford University, 1999.
- [21] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten. The weka data mining software: an update. *ACM SIGKDD explorations newsletter*,

- 11(1):10–18, 2009.
- [22] Philip N Howard and Bence Kollanyi. Bots, # strongerin, and # brexit: computational propaganda during the uk-eu referendum. *Available at SSRN 2798311*, 2016.
 - [23] Henry Hsu and Peter A Lachenbruch. Paired t test. *Wiley encyclopedia of clinical trials*, pages 1–3, 2007.
 - [24] Wayne Iba and Pat Langley. Induction of one-level decision trees. In *Machine Learning Proceedings 1992*, pages 233–240. Elsevier, 1992.
 - [25] Tobias R Keller and Ulrike Klinger. Social bots in election campaigns: Theoretical, empirical, and methodological implications. *Political Communication*, 36(1):171–189, 2019.
 - [26] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM, 2003.
 - [27] David MJ Lazer, Matthew A Baum, Yochai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, et al. The science of fake news. *Science*, 359(6380):1094–1096, 2018.
 - [28] Saskia Le Cessie and Johannes C Van Houwelingen. Ridge estimators in logistic regression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 41(1):191–201, 1992.
 - [29] S Li and ATS Ho. Hiding information in a digital environment. *Publication No. WO/2016/075459, International Application No. PCT/GB2015/053412, GB application filed on 11th November*, 2014.
 - [30] David J.C. Mackay. *Introduction to gaussian processes*, 1998.
 - [31] Silvia Mitter, Claudia Wagner, and Markus Strohmaier. A categorization scheme for socialbot attacks in online social networks. *arXiv preprint arXiv:1402.6288*, 2014.
 - [32] Jojo Moolayil, Moolayil, and Suresh John. *Learn Keras for Deep Neural Networks*. Springer, 2019.
 - [33] Atif Mushtaq, Todd Rosenberry, Ashar Aziz, and Ali Islam. Distributed systems and methods for automatically detecting unknown bots and botnets, February 5 2019. US Patent 10,200,384.
 - [34] Nic Newman, Richard Fletcher, Antonis Kalogeropoulos, David Levy, and Rasmus Kleis Nielsen. Reuters institute digital news report 2017. *Reuters Institute for the Study of Journalism*, 2017.
 - [35] Nic Newman, Richard Fletcher, Antonis Kalogeropoulos, David AL Levy, and Rasmus-Kleis Nielsen. Digital news report 2016. *Reuters Institute for the study of Journalism*, 2016.
 - [36] Sankar K Pal and Sushmita Mitra. Multilayer perceptron, fuzzy sets, and classification. *IEEE Transactions on neural networks*, 3(5):683–697, 1992.
 - [37] J. Ross Quinlan. Simplifying decision trees. *International Journal of Human-Computer Studies*, 27(3):221–234, 1987.
 - [38] Ross J. Quinlan. Learning with continuous classes. In *5th Australian Joint Conference on Artificial Intelligence*, pages 343–348, Singapore, 1992. World Scientific.
 - [39] Christopher Sandy, Patrice Rusconi, and Shujun Li. Can humans detect the authenticity of social media accounts? on the impact of verbal and non-verbal cues on credibility judgements of twitter profiles. In *2017 3rd IEEE International Conference on Cybernetics (CYBCONF)*, pages 1–8. IEEE, 2017.
 - [40] Saiph Savage, Andres Monroy-Hernandez, and Tobias Höllerer. Botivist: Calling volunteers to action using online bots. In *CSCW*. ACM, 2016.
 - [41] David Scuse and Peter Reutemann. Weka experimenter tutorial for version 3-5-5. *University of Waikato*, 2007.
 - [42] Tracy Jia Shen, Robert Cowell, Aditi Gupta, Thai Le, Amulya Yadav, and Dongwon Lee. How gullible are you?: Predicting susceptibility to fake news. In *Proceedings of the 11th ACM Conference on Web Science, WebSci 2019, Boston, MA, USA, June 30 - July 03, 2019*, pages 287–288, 2019.
 - [43] S.K. Shevade, S.S. Keerthi, C. Bhattacharyya, and K.R.K. Murthy. Improvements to the smo algorithm for svm regression. In *IEEE Transactions on Neural Networks*, 1999.

- [44] Pablo Suárez-Serrato, Margaret E Roberts, Clayton Davis, and Filippo Menczer. On the influence of social bots in online protests. In *International Conference on Social Informatics*, pages 269–278. Springer, 2016.
- [45] Onur Varol, Emilio Ferrara, Clayton A. Davis, Filippo Menczer, and Alessandro Flammini. Online human-bot interactions: Detection, estimation, and characterization. In *Proceedings of the Eleventh International Conference on Web and Social Media, ICWSM 2017, Montréal, Québec, Canada, May 15-18, 2017.*, pages 280–289, 2017.
- [46] Claudia Wagner, Silvia Mitter, Christian Körner, and Markus Strohmaier. When social bots attack: Modeling susceptibility of users in online social networks. In *# MSM*, pages 41–48, 2012.
- [47] Randall Wald, Taghi M Khoshgoftaar, Amri Napolitano, and Chris Sumner. Predicting susceptibility to social bots on twitter. In *2013 IEEE 14th International Conference on Information Reuse & Integration (IRI)*, pages 6–13. IEEE, 2013.
- [48] Cort J Willmott and Kenji Matsuura. Advantages of the mean absolute error (mae) over the root mean square error (rmse) in assessing average model performance. *Climate research*, 30(1):79–82, 2005.
- [49] Kai-Cheng Yang, Onur Varol, Clayton A Davis, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Technologies*, 1(1):48–61, 2019.
- [50] Xinyi Zhou and Reza Zafarani. Fake news: A survey of research, detection methods, and opportunities. *arXiv preprint arXiv:1812.00315*, 2018.

A Appendix

A.1 Feature sets

Name	Original Features	Features alteration
Botometer+ used in[7] From Botometer[45]	Macro categories scores: Sentiment, Friend, User, Content, Temporal, Net	(+) CAP (ENG and UNI) (+) Score (ENG and UNI) (+) #Tweets4WS (+) #Mentions4WS
ClassA- used in[7] From Fake Follower[12] (Class A)	friends/(followers2), age, #tweets, profile has name, has URL in profile, following rate, default image after 2 months, belongs to a list, profile has image, friends/followers \geq 50, ‘bot’ in biography, #friends, 2 \times followers \geq friends, #followers, friends/followers \simeq 100, no bio & no location & friends \geq 100, has address, has biography	(-) duplicated pictures of the profile

Table 5: Feature sets