2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)

Measuring the Information-Foraging Behaviors of Social Bots Through Word Usage

Zachary K. Stine, Tuja Khaund, and Nitin Agarwal
Department of Information Science
University of Arkansas at Little Rock
Little Rock, United States
{zkstine, txkhaund, nxagarwal}@ualr.edu

***Abstract*—Automated social bots are reported to account for a large sum of activity on social media sites such as Twitter. In this short paper, we study the information-foraging behaviors of social media users including bots. We present here a preliminary investigation which compares the behaviors of a set of suspected bots with non-automated accounts. To do so, we measure the distance between word distributions on a daily basis. We posit that this methodology provides a quantitative measure of behavior, which allows for more rigorous descriptions of bot behaviors that move beyond the assumption of bots as a monolithic category.**

Keywords—social bots, information foraging, social media, information theory

I. Introduction

Increasing attention has been given to the study of automated accounts on social media platforms, often referred to as bots, and their potential to meaningfully affect the social spaces in which they exist. While bot accounts may serve any number of purposes, benign or otherwise, much of the existing literature frames bots as a behaviorally monolithic category. This is reinforced by studies that differentiate solely between human and bot accounts. While such studies are important for distinguishing bots in general, there is a need to dig deeper into dynamic descriptions of bots in order to better understand the range of behaviors, and thus purposes, from which bots are programmed to act. To this end, we propose viewing the text produced by bot accounts as the record of an individual's information foraging behavior through a linguistic resource environment in order to generate textual content. The result of this foraging process may be text that is explorative or exploitative in nature. In this paper, we demonstrate one possible method for adapting a quantitative description of information-foraging behavior to describe text-producing agents. We apply this method to text generated over several days by a limited sample of Twitter users that includes suspected bots. We find that, within our limited sample, the suspected bots display behavior that is less explorative on average than the suspected human accounts.

II. Related Work

A. Social Bots

Reference [1] provides a review of some of the malicious roles played by bots in recent events as well as an overview of current bot detection methods. The Twitter-focused bot detection platform, *BotOrNot*, uses a Random Forest classifier to estimate the likelihood that an account is an automated bot [2][3][4]. Notably, the temporal features described in [3] are limited to users' tweet rates and time intervals between tweets. Similar features as well as sentiment variance and average number of tweets per user are described as classification features in [4]. While these features certainly capture important aspects of bot behaviors, they do not take into account how explorative or exploitative a bot's word usage is. By measuring the information-foraging behavior of accounts, we provide a first step toward richer behavioral descriptions and bot taxonomy.

B. Information Foraging

The rationale for the methodology we employ is adapted from [5], which studies the decision-making process of Charles Darwin in terms of the books he chose to read during key parts of his life. The authors view Darwin as an information forager selecting among books to read. They measure the Kullback-Leibler Divergence (KLD) across the topics from Darwin's reading list, inferred using Latent Dirichlet Allocation (LDA). The authors interpret the KLD between two topic distributions as a measure of surprise, which is measured in two ways: first, as the surprise from book-to-book (local surprise) and second, as the surprise from all books previously read to the next single book (global surprise). For additional background on information foraging, we encourage readers to see sources cited within [5].

An in-depth explanation of information theoretic quantities is provided by [6], including descriptions of KLD, used by [5], and Jensen-Shannon Divergence (JSD), used in this study. Of particular note is that KLD is not finite in cases where one distribution contains zero for an element that is nonzero in the distribution being compared. Here, we use JSD, which is always finite and therefore convenient for comparing word usage distributions which may be sparse. However, the choice of JSD alters the interpretation of information foraging as related to KLD in [5]. For our purposes, JSD can be simply understood as the distance between two distributions.

III. Methods

In this preliminary study, we analyze text from 32,641 tweets generated by a sample of 72 Twitter users between March 11 through April 17 of 2018. Of these 72 users, 32 are

suspected to be bot accounts while the remaining 40 accounts are suspected to be operated directly by human users. Our categorization of users as suspected bots or humans is based on *BotOrNot* scores where suspected bots have a minimum score of 0.79 and suspected humans have a maximum score of 0.05. Within this data set, the average number of tweets per user is 699.3 tweets for bots, and 256.6 tweets for humans. The average number of days on which an account tweeted is 24.3 for bots and 26.1 for humans. For each account, we tokenize the text from their tweets and compute term frequencies for each day the account tweeted on. We then compute the JSD between each day—the local surprise—and between all previous days and then next day—the global surprise.

IV. Results

In order to make a high level comparison between the suspected bot and human accounts, we examine the average local and global JSD values for each user. We find that, on average, suspected bots have lower values of local and global JSD and thus are less linguistically explorative than the suspected human accounts (Fig. 1). Additionally, we find that the sampled bots have greater variance in the average local and global JSD values than the sampled humans.

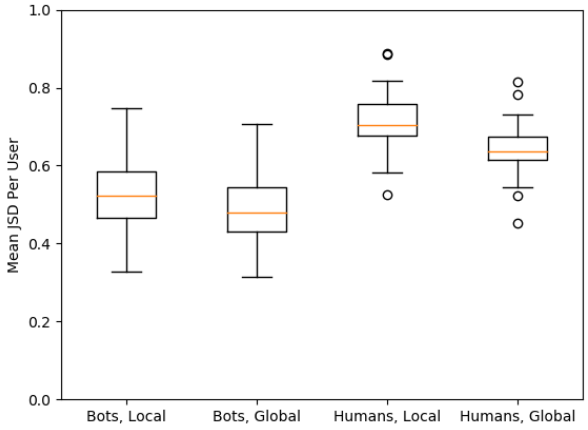

Fig. 1. From left to right, the mean local and global JSD values for bots, followed by the same for humans. The sampled humans show higher JSD values on average, implying more explorative word usage than the sampled bots.

V. Discussion

This study represents a first step toward richer behavioral descriptions of social media users in general, and bots in particular. Instead of distinguishing solely between bots and humans, we can add additional complexity to these distinctions by now talking about linguistically exploitative bots, explorative bots, exploitative humans, or explorative humans. If a user writes about similar things from day to day, then we should see low surprise in both local and global measures. On the other hand, if a user writes about a new set of topics every day, then both global and local surprise should be continually high. However, if a user switches back and forth between two sets of distinct word usage patterns from day to day, then the local surprise will stay high, but the global surprise will become low. Thus, bots that have a focused purpose may have a focused vocabulary, resulting in exploitative foraging behavior. Yet bots whose purposes do not include a single focused message may show higher levels of surprise in their word usage. By measuring how explorative or exploitative a bot's text production is, we can infer more about the intent of the bot by measuring how focused the bot's messages are over time.

Due in part to the preliminary nature of this study, there are several limitations that will need to be overcome in future work as we explore this method more deeply. First, a larger sample of users will be necessary to make broader generalizations about the information-foraging behaviors of users. Broader generalizations will also become more robust as we expand the time range of tweets collected from around a month to longer durations. Additionally, the features analyzed in this study, token frequencies, are quite fine-grained and thus liable to fluctuate greatly. In order to minimize the sensitivity of these features, we look at the word usage over a given day, but there are likely other ways of coarse-graining the text that should be investigated as well (*e.g.* the topic modeling approach taken in [5]).

As the sophistication of bots increases, the sophistication of our descriptions must increase as well. We have shown that within a sample of users, suspected bots are less explorative in their word usage than suspected humans. We have also observed greater variance in the information-foraging behavior of the suspected bots than in the humans. While these results are preliminary, they demonstrate the types of questions that can be answered using this method.

Acknowledgments

This research is funded in part by the U.S. National Science Foundation (IIS-1636933, ACI-1429160, and IIS-1110868), U.S. Office of Naval Research (N00014-10-1-0091, N00014-14-1-0489, N00014-15-P-1187, N00014-16-1-2016, N00014-16-1-2412, N00014-17-1-2605, N00014-17-1-2675), U.S. Air Force Research Lab, U.S. Army Research Office (W911NF-16-1-0189), U.S. Defense Advanced Research Projects Agency (W31P4Q-17-C-0059) and the Jerry L. Maulden/Entergy Endowment at the University of Arkansas at Little Rock.

References

[1] E. Ferrara et al., "The rise of social bots," Commun. ACM, vol. 59(7), pp. 96–104, June 2016. DOI: https://doi.org/10.1145/2818717

[2] C. Davis et al., "BotOrNot: a sytem to evaluate social bots," in Proc. of the 25th International World Wide Web Conference, pp. 273–274, 2016.

[3] O. Varol et al., "Online human-bot interactions: detection, estimation, and characterization," International AAAI Conference on Web and Social Media, pp. 280–289, March, 2017.

[4] V. Subrahmainian et al., "The DARPA Twitter bot challenge," Computer vol. 59(6), pp. 38-46, June 2016. DOI: 10.1109/MC.2016.183

[5] J. Murdock, C. Allen, and S. DeDeo, "Exploration and exploitation of Victorian science in Darwin's reading notebooks," Cognition, vol. 159, pp. 117–126, February 2017.

[6] S. DeDeo et al., "Bootstrap methods for the empirical study of decision-making and information flows in social systems," Entropy, vol. 15(6), pp. 2246–2276, June 2013. DOI: 10.3390/e15062246