Taming Social Bots: Detection, Exploration and Measurement

Abdullah Mueen
The University of New Mexico
mueen@unm.edu

Nikan Chavoshi
Oracle Corporation
nikan.chavoshi@gmail.com

Amanda Minnich
Lawrence Livermore Nat'l Lab
amandajean119@gmail.com

ABSTRACT

Social bots have been around since 2008, and thus, they have been polluting our online spaces for over a decade. Social bots are capable of swaying political opinion, spreading false information, and recruiting for terrorist organizations. Social bots use various sophisticated techniques by adopting emotions, sympathy following, synchronous deletions, and profile molting.

There are several approaches proposed in the literature for detecting, exploring, and measuring of social bots. We provide a comprehensive overview of the existing work from the data mining and machine learning perspective, discuss relative strengths and weaknesses of various methods, make recommendations for researchers and practitioners, and propose novel directions for future research in taming social bots. This tutorial also discusses pitfalls in collecting and sharing data on social bots.

CCS CONCEPTS

• Information systems → Web and social media search;

KEYWORDS

Social bots, campaign, purge, link farming

1 TOPIC OVERVIEW

The massive changes in our lives wrought by social media are well-known. Thanks to services like Twitter, Facebook, and Instagram, we can now interact in ways that would have sounded like science fiction thirty years ago. With these new platforms come new ways to abuse them. From viral marketing schemes to spreading fake news (e.g. the Pizzagate conspiracy theory), social media is full of people looking to misuse the connections it creates.

Bots, automated accounts run by computer programs, are a noteworthy way of abusing social media. Bots are quick to construct, almost free, and increasingly common [10]. Based on our research, a Twitter bot can be constructed with about 30 seconds of work and requires negligible resources to run. There are as many as 49.2 million bots on Twitter [22], a number that grows continuously. Although they can have benign uses, bots are also used to spread false information [16], pump up stock prices [8], harass the opponents of authoritarian regimes, and recruit for terrorist organizations [14].

There are currently research centered around detecting bots [24][15][20][6]. Technologies like DeBot [4] and BotOrNot [9] are

good at finding certain types, with DeBot specializing in detecting swarms of synchronous bots [3] and BotOrNot specializing in English language bots. There has been research on exploration techniques to improve the recall rate of initial detection systems [18][19]. There has been research on understanding bot behavior [12][23][5][13][17] and their impact on social and political events under various geographical [21][11][2], cultural, financial [8], behavioral [12] and social [7] contexts as well. Tutorial slides are available here [1].

2 COVERAGE

The topics covered in this tutorial have been published in KDD, ICDM, WWW, ASONAM, ICWSM and WSDM conferences in the past. Several of the papers we cover in this tutorial have won best paper awards, and received large number of citations in a short time. Data mining conferences always host at least one dedicated session on security issues in social media and several sessions on social media mining, in general. CIKM has had tutorials on social media mining, for example “Malware Analysis for Data Scientists” presented by Charles Nicholas in CIKM 2017. However, there has not been a tutorial on social bots, despite their prevalence. We think the following three categories of people will be interested.

Data mining/Database researchers: Researchers and those working on specific social media mining problems will find the tutorial informative. We will have a collection of pointers to publicly available datasets and online tools. Also, as the tutorial ends with a discussion of open problems to work on in the area, graduate students looking for an interesting problem in a hot area will be well-served.

Data mining/Database educators: Professors who teach cybersecurity and data mining can gain benefit from this tutorial's slides. Such individuals will receive our comprehensive (and modifiable) slides and observe our presentation of them. They will be able to base 10 to 20 hours of graduate instruction on the tutorial.

Data mining/Database application developers: These developers will learn the latest techniques for mining streaming data (e.g. tweets) and see examples of how the techniques fit into real-life applications (e.g. bot detection). They will also find the collection of code readable and useful once the algorithms are understood.

3 PRESENTER BIOS

Dr. Abdullah Mueen is an Associate Professor in Computer Science at University of New Mexico. Earlier he was a Scientist in the Cloud and Information Science Lab at Microsoft Corporation. His major interest is in temporal data mining with a focus on two unique types of signals: social networks and electrical sensors. He has been actively publishing in the data mining conferences including KDD, ICDM and SDM, and journals including DMKD and KAIS. He has received runner-up award in the Doctoral Dissertation Contest in KDD 2012. He has won the best paper award in

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CIKM '19, November 3–7, 2019, Beijing, China
© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-6976-3/19/11...\$15.00
<https://doi.org/10.1145/3357384.3360315>

SIGKDD 2012. His research is funded by NSF, DARPA and AFRL. Earlier, he earned PhD degree at the University of California at Riverside and BSc degree at Bangladesh University of Engineering and Technology.

Dr. Nikan Chavoshi has joined Oracle as a Senior Member of Technical Staff in June 2018. Earlier, she earned her PhD degree in Computer Science from the University of New Mexico. Her research interest is in time series mining and temporal activity analysis. Her graduate work was on analyzing temporal behavior of automated accounts in Twitter, for which, she was awarded the Outstanding Graduate Student for the CS Department in 2018. In her PhD, she worked with Dr. Abdullah Mueen and designed a near real-time system, named DeBot, to detect automated accounts in Twitter. Dr. Chavoshi has published research articles in top web mining venues including WWW, ICDM, SocInfo, ASONAM and KAIS.

Dr. Amanda J. Minnich is a Machine Learning Research Scientist and Molecular Data-Driven Modeling Team Lead at Lawrence Livermore National Lab (LLNL). At LLNL she is part of the multi-institution ATOM Consortium, where she applies Machine Learning techniques to biological data for drug discovery purposes. Dr. Minnich received a BA in Integrative Biology from UC Berkeley (2009) and an MS (2014) and PhD with Distinction (2017) in Computer Science from the University of New Mexico. While at UNM she was named an NSF Graduate Research Fellow, a PiBBs Fellow, a Grace Hopper Scholar, and the Outstanding Graduate Student for the CS Department in 2017. She has published her work at and served on Program Committees for top conferences including WWW, ASONAM, KDD, ICDM, SC, GTC, and ICWE, and has been issued a patent for her dissertation work. Dr. Minnich also has a passion for advocating for women in tech; she co-founded and served as President of UNM's first chartered Women in Computing group, she frequently volunteers at women in tech events, and she will be co-chairing the Artificial Intelligence track at Grace Hopper Celebration 2019.

4 ACKNOWLEDGMENT

The material is based on work supported by the National Science Foundation Grant No. SHF-1527127 and OIA-1757207.

REFERENCES

- [1] [n. d.]. Tutorial Page: ([n. d.]). <https://www.cs.unm.edu/~mueen/Tutorial/TamingBots.html>.
- [2] Noor Abu-El-Rub and Abdullah Mueen. 2019. BotCamp: Bot-driven Interactions in Social Campaigns. In *The Web Conference*.
- [3] Alex Beutel, Wanhong Xu, Venkatesan Guruswami, Christopher Palow, and Christos Faloutsos. 2013. CopyCatch: Stopping group attacks by spotting lockstep behavior in social networks. In *Proceedings of the 22nd international conference on World Wide Web - WWW '13*. ACM Press, New York, New York, USA, 119–130. <https://doi.org/10.1145/2488388.2488400>
- [4] Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. 2016. DeBot: Twitter Bot Detection via Warped Correlation. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 817–822. <https://doi.org/10.1109/ICDM.2016.0096>
- [5] Nikan Chavoshi and Abdullah Mueen. 2018. Model Bots, not Humans on Social Media. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 178–185. <https://doi.org/10.1109/ASONAM.2018.8508279>
- [6] Alceu Ferraz Costa, Yuto Yamaguchi, Agma Juci Machado Traina, Caetano Traina Jr., and Christos Faloutsos. 2015. RSC: Mining and Modeling Temporal Activity in Social Media. In *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '15)*. ACM.
- [7] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2017. Social fingerprinting: detection of spambot groups through DNA-inspired behavioral modeling. *IEEE Transactions on Dependable and Secure Computing* (2017), 1–1. <https://doi.org/10.1109/TDSC.2017.2681672>
- [8] Stefano Cresci, Fabrizio Lillo, Daniele Regoli, Serena Tardelli, and Maurizio Tesconi. 2019. Cashtag Piggybacking: Uncovering Spam and Bot Activity in Stock Microblogs on Twitter. *ACM Transactions on the Web* 13, 2 (4 2019), 1–27. <https://doi.org/10.1145/3313184>
- [9] Clayton Allen Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. BotOrNot: A System to Evaluate Social Bots. In *Proceedings of the 25th International Conference Companion on World Wide Web - WWW '16 Companion*. ACM Press, New York, New York, USA, 273–274. <https://doi.org/10.1145/2872518.2889302>
- [10] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2014. The Rise of Social Bots. *CoRR abs/1407.5* (2014). <http://arxiv.org/abs/1407.5225>
- [11] Michelle C Forelle, Philip N. Howard, Andres Monroy-Hernandez, and Saiph Savage. 2015. Political Bots and the Manipulation of Public Opinion in Venezuela. *SSRN Electronic Journal* (2015). <https://doi.org/10.2139/ssrn.2635800>
- [12] Andrew Hall, Loren Terveen, and Aaron Halfaker. 2018. Bot Detection in Wikidata Using Behavioral and Other Informal Cues. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (11 2018), 1–18. <https://doi.org/10.1145/3274333>
- [13] Hossein Hamooni, Nikan Chavoshi, and Abdullah Mueen. 2016. On URL Changes and Handovers in Social Media. Springer International Publishing, 58–74. https://doi.org/10.1007/978-3-319-47880-7_4
- [14] Jytte Klausen, Christopher Marks, and Tauhid Zaman. 2016. Finding Online Extremists in Social Networks. (10 2016). <http://arxiv.org/abs/1610.06242>
- [15] Thai Le, Kai Shu, Maria D Molina, Dongwon Lee, S. Shyam Sundar, and Huan Liu. 2019. 5 Sources of Clickbaits You Should Know! Using Synthetic Clickbaits to Improve Prediction and Distinguish between Bot-Generated and Human-Written Headlines. In *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. 33–40. <https://doi.org/10.1145/3341161.3342875>
- [16] Kyumin Lee, James Caverlee, Zhiyuan Cheng, and Daniel Z. Sui. 2011. Content-driven detection of campaigns in social media. In *Proceedings of the 20th ACM international conference on Information and knowledge management - CIKM '11*. ACM Press, New York, New York, USA, 551. <https://doi.org/10.1145/2063576.2063658>
- [17] Enrico Mariconti, Jeremiah Onaolapo, Syed Sharique Ahmad, Nicolas Nikiforou, Manuel Egele, Nick Nikiforakis, and Gianluca Stringhini. 2017. What's in a Name?: Understanding Profile Name Reuse on Twitter. In *Proceedings of the 26th International Conference on World Wide Web - WWW '17*. ACM Press, New York, New York, USA, 1161–1170. <https://doi.org/10.1145/3038912.3052589>
- [18] A. Minnich, N. Chavoshi, D. Koutra, and A. Mueen. 2017. BotWalk: Efficient adaptive exploration of twitter bot networks. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2017*. <https://doi.org/10.1145/3110025.3110163>
- [19] F Morstatter, K M Carley, and H Liu. 2015. Bot Detection in Social Media: Networks, Behavior, and Evaluation. In *ASONAM - Tutorial*.
- [20] Pujan Paudel, Trung T. Nguyen, Amartya Hatua, and Andrew H. Sung. 2019. How the Tables Have Turned: Studying the New Wave of Social Bots on Twitter Using Complex Network Analysis Techniques. In *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. 501–508. <https://doi.org/10.1145/3341161.3342898>
- [21] Pablo Suárez-Serrato, Margaret E. Roberts, Clayton Davis, and Filippo Menczer. 2016. On the Influence of Social Bots in Online Protests. Springer, Cham, 269–278. https://doi.org/10.1007/978-3-319-47874-6_19
- [22] V. S. Subrahmanian, Amos Azaria, Skylar Durst, Vadim Kagan, Aram Galstyan, Kristina Lerman, Linhong Zhu, Emilio Ferrara, Alessandro Flammini, Filippo Menczer, Rand Waltzman, Andrew Stevens, Alexander Dekhtyar, Shuyang Gao, Tad Hogg, Farshad Kooti, Yan Liu, Onur Varol, Prashant Shiralkar, Vinod Vydiswaran, Qiaozhu Mei, and Tim Huang. 2016. The DARPA Twitter Bot Challenge. (1 2016). <http://arxiv.org/abs/1601.05140>
- [23] Kurt Thomas, Chris Grier, Dawn Song, and Vern Paxson. 2011. Suspended accounts in retrospect: an analysis of twitter spam. In *Proceedings of the 2011 ACM ... (IMC '11)*. 243–258. <https://doi.org/10.1145/2068816.2068840>
- [24] Mengfan Yao, Charalampos Chelmis, and Daphney-Stavroula Zois. 2019. Cyberbullying Ends Here: Towards Robust Detection of Cyberbullying in Social Media. In *The World Wide Web Conference on - WWW '19*. ACM Press, New York, New York, USA, 3427–3433. <https://doi.org/10.1145/3308558.3313462>