

340k unique users using a generic methodology (which can also be used for other platforms and other case studies). Although focused on a narrow slice of time, we found that, in general, users tweeting about Gamergate appear to be Twitter savvy and quite engaged with the platform. They produce more tweets than random users, and have more friends and followers as well. Surprisingly, we observed that, while expressing more negative sentiment overall, these users only differed significantly from random users with respect to joy. Finally, we looked at account suspension, finding that Gamergate users are less likely to be suspended due to the inherent difficulties in detecting and combating online harassment activities.

While we believe our work contributes to understanding large-scale online harassment, it is only a start. As part of future work, we plan to perform a more in-depth study of Gamergate, focusing on how it evolved over time. Overall, we argue that a deeper understanding of how online harassment campaigns function can enable our community to better address them and propose detection tools as well as mitigation strategies.

Acknowledgements. This work has been funded by the European Commission as part of the ENCASE project (H2020-MSCA-RISE), under GA number 691025.

6. REFERENCES

- [1] Anonymous. "i dated zoe quinn". 4chan. <https://archive.is/qrS5Q>.
- [2] Anonymous. Zoe Quinn, prominent SJW and indie developer is a liar and a slut. 4chan. <https://archive.is/QIjm3>.
- [3] D. Chatzakou, N. Kourtellis, J. Blackburn, E. D. Cristofaro, G. Stringhini, and A. Vakali. Mean Birds: Detecting Aggression and Bullying on Twitter. *arXiv preprint, arXiv:1702.06877*, 2017.
- [4] D. Chatzakou, V. A. Koutsonikola, A. Vakali, and K. Kafetsios. Micro-blogging content analysis via emotionally-driven clustering. In *ACII*, 2013.
- [5] C. Chen, J. Zhang, X. Chen, Y. Xiang, and W. Zhou. 6 million spam tweets: A large ground truth for timely Twitter spam detection. In *IEEE ICC*, 2015.
- [6] Y. Chen, Y. Zhou, S. Zhu, and H. Xu. Detecting Offensive Language in Social Media to Protect Adolescent Online Safety. In *PASSAT and SocialCom*, 2012.
- [7] K. Dinakar, R. Reichart, and H. Lieberman. Modeling the detection of textual cyberbullying. *The Social Mobile Web*, 11, 2011.
- [8] N. Djuric, J. Zhou, R. Morris, M. Grbovic, V. Radosavljevic, and N. Bhamidipati. Hate Speech Detection with Comment Embeddings. In *WWW*, 2015.
- [9] M. Giatoglou, D. Chatzakou, N. Shah, A. Beutel, C. Faloutsos, and A. Vakali. Nd-sync: Detecting synchronized fraud activities. In *Advances in Knowledge Discovery and Data Mining, 19th Pacific-Asia Conference*, 2015.
- [10] A. Go, R. Bhayani, and L. Huang. Twitter sentiment classification using distant supervision. *Processing*, 2009.
- [11] J. Guberman and L. Hemphill. Challenges in modifying existing scales for detecting harassment in individual tweets. In *System Sciences*, 2017.
- [12] A. Hern. Feminist critics of video games facing threats in 'gamergate' campaign. *The Guardian*, Oct 2014. <https://www.theguardian.com/technology/2014/oct/23/felicia-days-public-details-online-gamergate>.
- [13] G. E. Hine, J. Onaolapo, E. De Cristofaro, N. Kourtellis, I. Leontiadis, R. Samaras, G. Stringhini, and J. Blackburn. A longitudinal measurement study of 4chan's politically incorrect forum and its effect on the web. *arXiv preprint arXiv:1610.03452*, 2016.
- [14] H. Hosseinmardi, R. Han, Q. Lv, S. Mishra, and A. Ghasemianlangroodi. Towards understanding cyberbullying behavior in a semi-anonymous social network. In *IEEE/ACM ASONAM*, 2014.
- [15] H. Hosseinmardi, S. A. Mattson, R. I. Rafiq, R. Han, Q. Lv, and S. Mishra. Analyzing Labeled Cyberbullying Incidents on the Instagram Social Network. In *SocInfo*, 2015.
- [16] I. Kayes, N. Kourtellis, D. Quercia, A. Iamnitchi, and F. Bonchi. The Social World of Content Abusers in Community Question Answering. In *WWW*, 2015.
- [17] A. Massanari. #gamergate and the fapping: How reddit's algorithm, governance, and culture support toxic technocultures. *New Media & Society*, 2015.
- [18] T. E. Mortensen. Anger, fear, and games. *Games and Culture*, 2016.
- [19] G. Navarro. A Guided Tour to Approximate String Matching. *ACM Comput. Surv.*, 33(1), 2001.
- [20] C. Nobata, J. Tetreault, A. Thomas, Y. Mehdad, and Y. Chang. Abusive Language Detection in Online User Content. In *WWW*, 2016.
- [21] Pham, Sherisse. Twitter tries new measures in crackdown on harassment. *CNNtech*, February 2017. <http://money.cnn.com/2017/02/07/technology/twitter-combat-harassment-features/>.
- [22] Twitter trolls are now abusing the company's bottom line. <http://www.salon.com/2016/10/19/twitter-trolls-are-now-abusing-the-companys-bottom-line/>, 2016.
- [23] H. Sanchez and S. Kumar. Twitter bullying detection. *ser. NSDI*, 12, 2011.
- [24] G. Stringhini, C. Kruegel, and G. Vigna. Detecting spammers on social networks. In *ACSAC*, 2010.
- [25] A. Sulleyman. Twitter temporarily limiting users for abusive behaviour. *Independent*, February 2017. <https://www.theguardian.com/technology/2014/oct/23/felicia-days-public-details-online-gamergate>.
- [26] The Guardian. Did trolls cost Twitter 3.5bn and its sale? <https://www.theguardian.com/technology/2016/oct/18/did-trolls-cost-twitter-35bn>, 2016.
- [27] C. Van Hee, E. Lefever, B. Verhoeven, J. Mennes, B. Desmet, G. De Pauw, W. Daelemans, and V. Hoste. Automatic detection and prevention of cyberbullying. In *Human and Social Analytics*, 2015.
- [28] A. H. Wang. Don't follow me: Spam detection in Twitter. In *SECRYPT*, 2010.
- [29] N. Wingfield. Feminist critics of video games facing threats in 'gamergate' campaign. *New York Times*, Oct 2014. <https://www.nytimes.com/2014/10/16/technology/gamergate-women-video-game-threats-anita-sarkeesian.html>.