BOTS AMONG US: PREVALENCE, INFLUENCE, AND ROLES OF AUTOMATED ACCOUNTS IN THE GERMAN TWITTER FOLLOW NETWORK

Felix Victor Münch
Leibniz Institute for Media Research | Hans-Bredow-Institut (HBI)

Cornelius Puschmann
Leibniz Institute for Media Research | Hans-Bredow-Institut (HBI)

Ben Thies
Leibniz Institute for Media Research | Hans-Bredow-Institut (HBI)

Axel Bruns
Queensland University of Technology

Background

Social bots are undermining trust in social media. They spread low-credibility content (Shao et al., 2018), so-called fake news (Vosoughi, Roy, & Aral, 2018), and spam (Bruns et al., 2018). However, most research analyses data based on the active sharing of links, keywords, or hashtags rather than assessing the longer-term presence of bots as an integral part of platforms.

To address this gap, we present what to our knowledge is the first study that assesses the prevalence, influence, and roles of automated accounts in a Twitter follow network on a national scale: the German-speaking Twittersphere. This work in progress allows us to analyse the long-term structural role, impact, and possible audience of bots beyond the context of single events and topics.

Method

To collect a follow network of the most central accounts in the German-speaking Twittersphere, we use the rank-degree method, a graph exploration method that is able to identify the most influential spreaders within complex networks (Salamanos, Voudigari, & Yannakoudakis, 2017). As this walk-based technique only requires node-local information, we were able to adapt it as a data mining method using the cost-free standard Twitter API only.
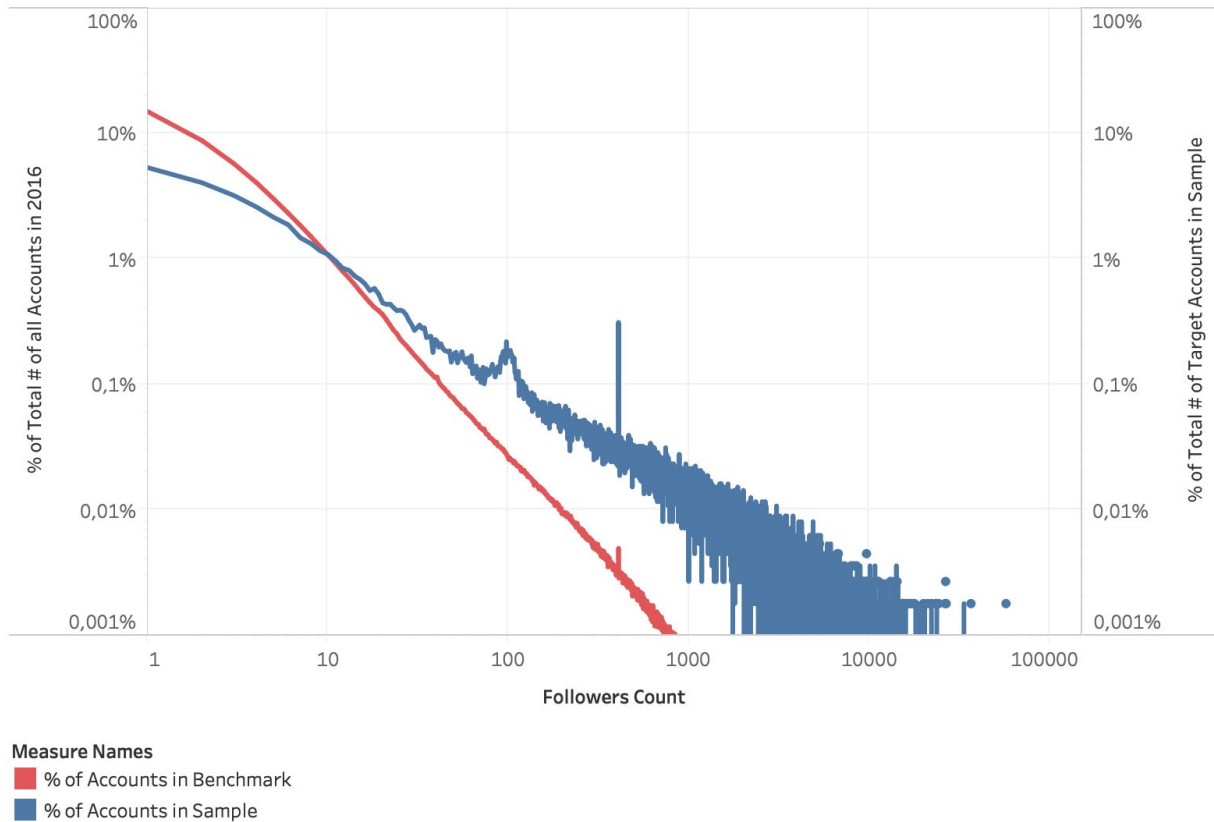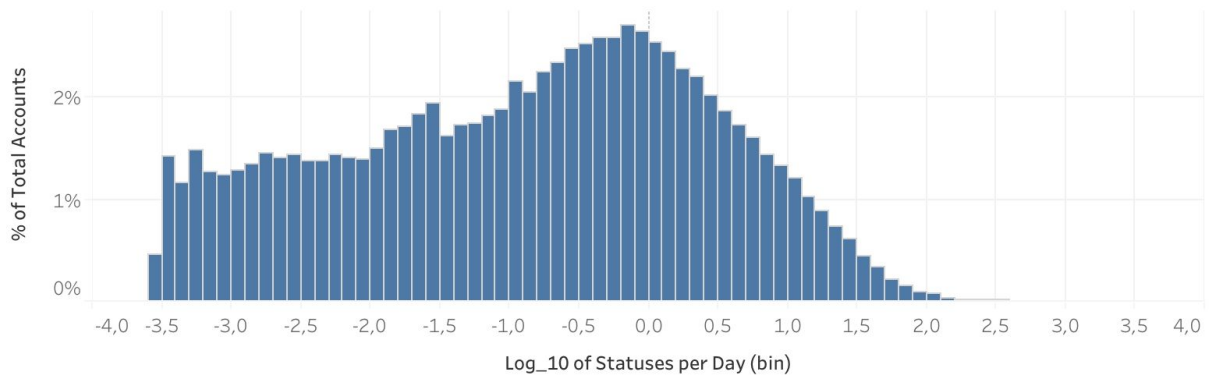

Figure 1: Comparison of the distribution of follower counts in our sample (blue) and the complete German account data from 2016 (red).

So far, we have compiled a network of more than 200,000 accounts and 840,000 edges. A detailed description of our collection method and more insights into the sample quality will be part of forthcoming publications.

A preliminary comparison of our sample with a benchmark dataset (comprising the details of all Twitter accounts who had set their interface language to German in 2016) shows that our sample contains an influential sub-population with large (Figure 1), recently active (Figures 2 & 3) audiences.

As shown by Münch (2019), a sample of the most followed accounts in the Australian Twittersphere sufficed to identify densely connected topical clusters, overall reflecting results from the full dataset. Therefore, we assume that our sample allows to identify long-lasting topical sub-publics which reflect the macro-structure of the German Twittersphere (Figure 5).
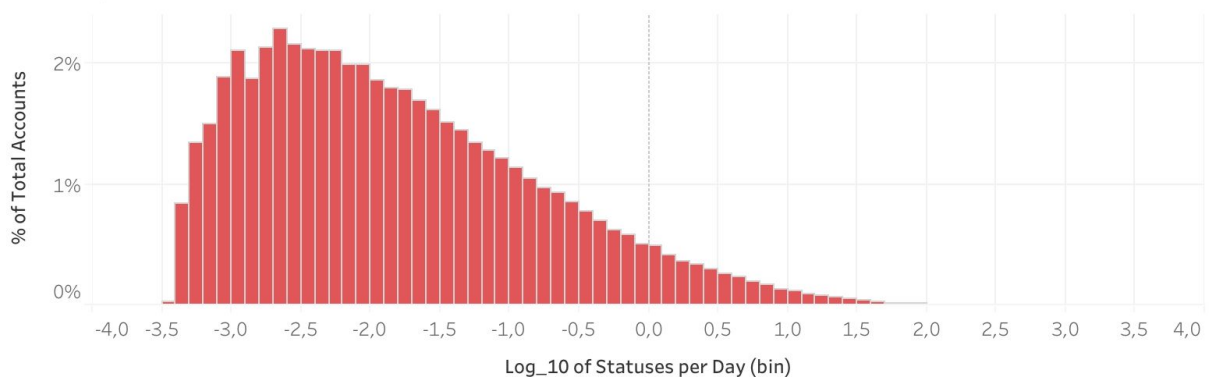

Figure 2: Comparison of the activity distribution of statuses per day in our sample (blue) and the complete German account data from 2016 (red).

To identify bots, we employ the Botometer API (https://botometer.iuni.iu.edu/), which returns a score between 0 and 1 that indicates the probability of an account to be automated. Botometer is based on Bayesian statistics and a set of machine learning approaches that take over 1,000 features into account (Varol, Ferrara, Davis, Menczer, & Flammini 2017). It represents one of the most widely used methods to identify bots in academic research today.

Our network sample and the automation scores allow us to localise bots within topical clusters, to assess their potential influence (by in-degree, page rank, harmonic closeness centrality, betweenness centrality, and k-coreness), and to qualitatively assess the role of the most central bots.
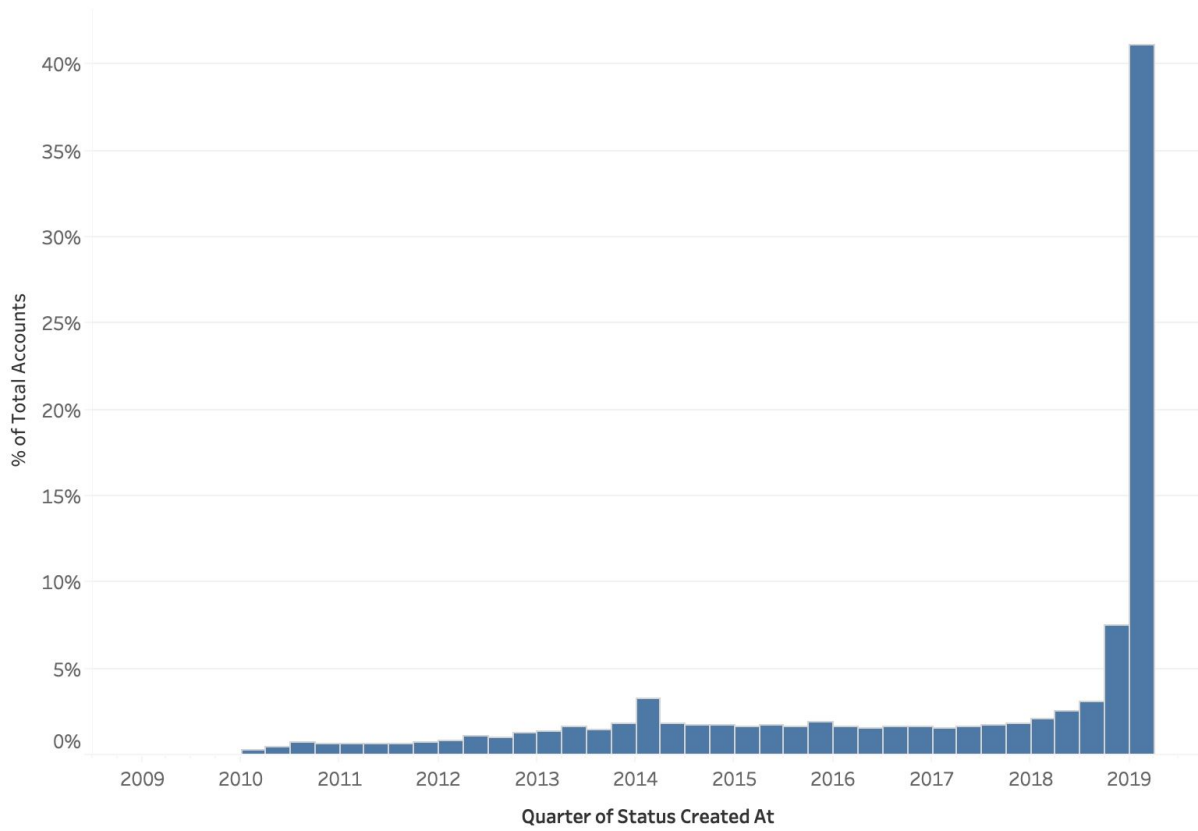
Last Status by Accounts in Sample

Figure 3: Distribution of the date of the last status by accounts in our sample.

Preliminary Findings

As the data collection and classification is ongoing, we can only discuss preliminary findings. When presenting this paper, we will have concluded our analysis.

Overall, using an automation score threshold of 0.75, we find that 11% of all accounts in the data classified so far (n = 91,000) meet this criterion (Figure 4). All of these accounts have a comparably low page rank and in-degree (Figures 6 & 7).

Most bots are concentrated in a cluster dominated by PR and marketing accounts (red in Figure 5) or in isolated bot-nets (Figures 6 & 7). The cluster dominated by political and news accounts (purple/black in Figure 5) remains largely unaffected.

Qualitatively examining central bots we find that many of them are benign and/or inactive. For example, the *verified* account of the German a-capella band Wise Guys (https://twitter.com/wise_guys) was automated, but only shared their blog posts until they split up in 2017.
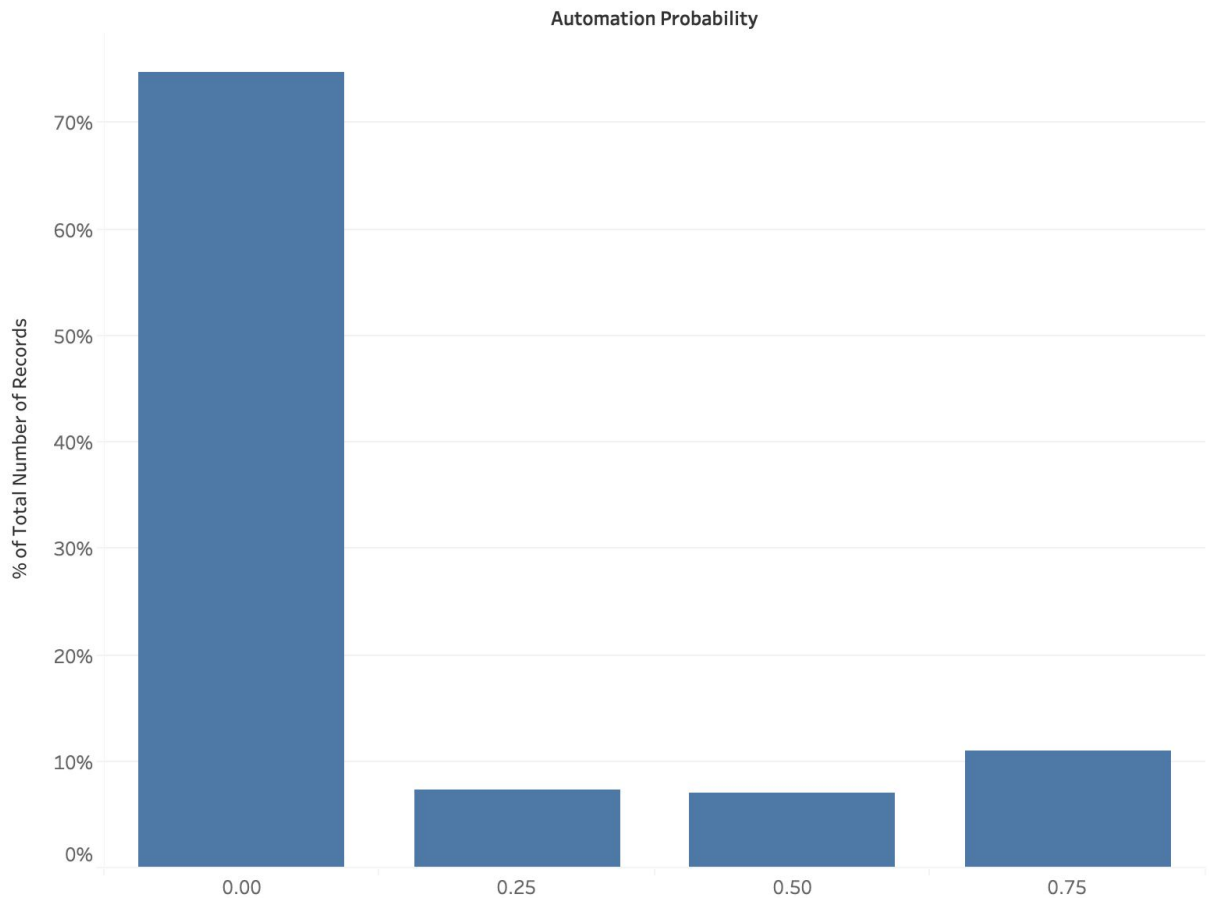
Figure 4: Distribution of automation probability of accounts as determined by Botometer.

Conclusion

Within the influential population of our sample, automated accounts are prevalent at a percentage that confirms estimates of other analyses. Varol et al. (2017), for example, estimate that 9% to 15% of the active Twitter population are automated.

However, based on their low centrality and cluster locations, bots seem to have a low potential to spread information outside of commercial contexts. Spot testing confirms that inactivity leads to high Botometer scores. Our estimate of bot prevalence will likely drop when we account for this. Also, not all automation is malignant, but, for example, helps content creators to serve several platforms at once.

This indicates that bots have only a low negative impact on the German-speaking Twittersphere. However, the most sophisticated bots will likely remain absent from our study, as false negatives. Trolls and semi-automated accounts necessitate further research.
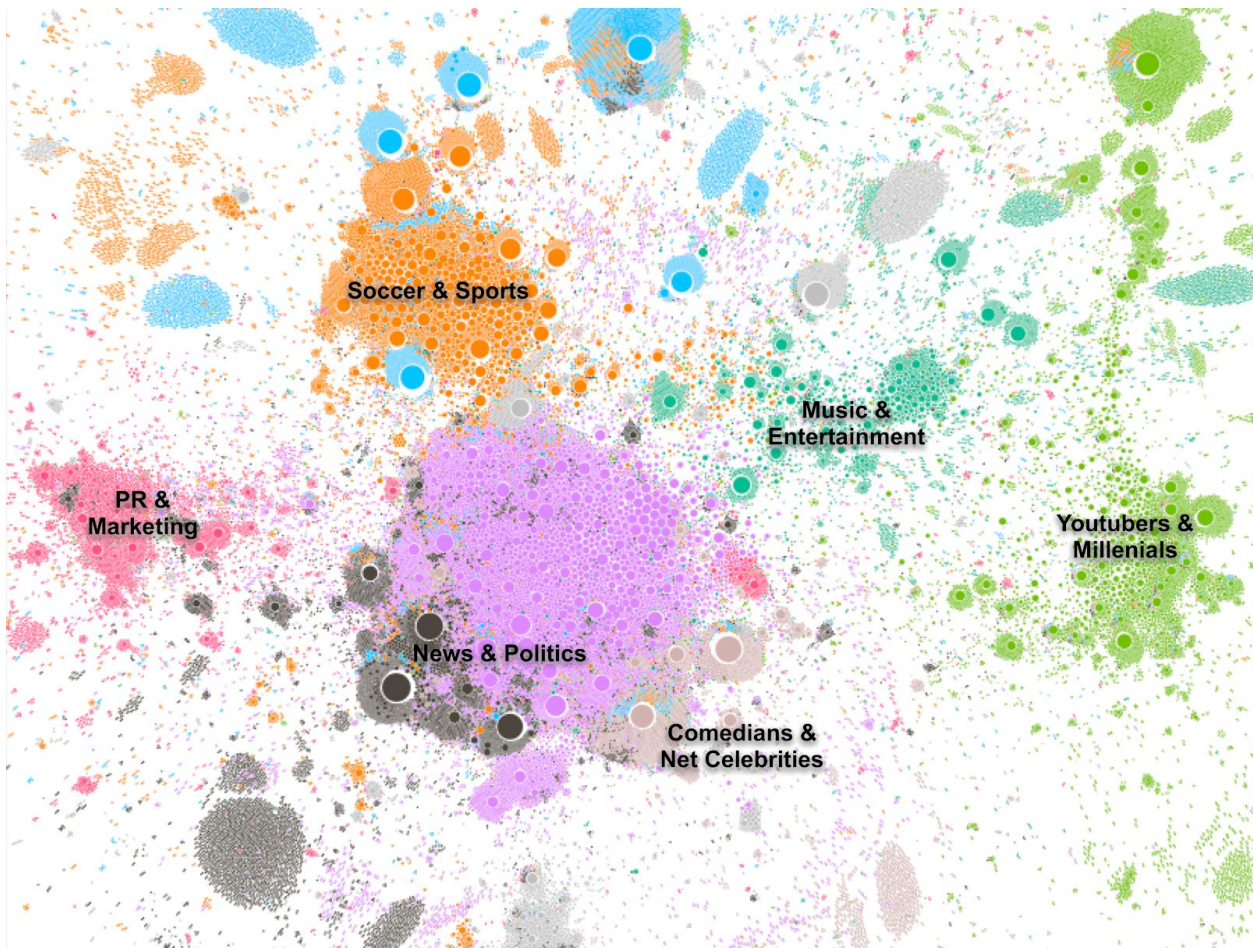
Figure 5: Detail of a network visualisation of the German Twittersphere sample with topical sub-clusters of nodes, coloured by selected groups identified via modularity maximisation. Size of nodes is proportional to PageRank. Links are hidden. (Layout based on Force-Atlas 2 with Lin-Log-Mode and No-Overlap applied as implemented in Gephi 0.9.2.)

Outlook

Our new sampling approach for large-scale follow networks combined with the bot-detection API of Botometer is also promising in other contexts: for example, Twitterspheres based on other languages. We will publish the code for collecting the data necessary for this analysis under an open source license. For this dataset, our approach opens further avenues of enquiry, such as the long-term monitoring of automated accounts in the German Twittersphere, including content analyses of their activities.
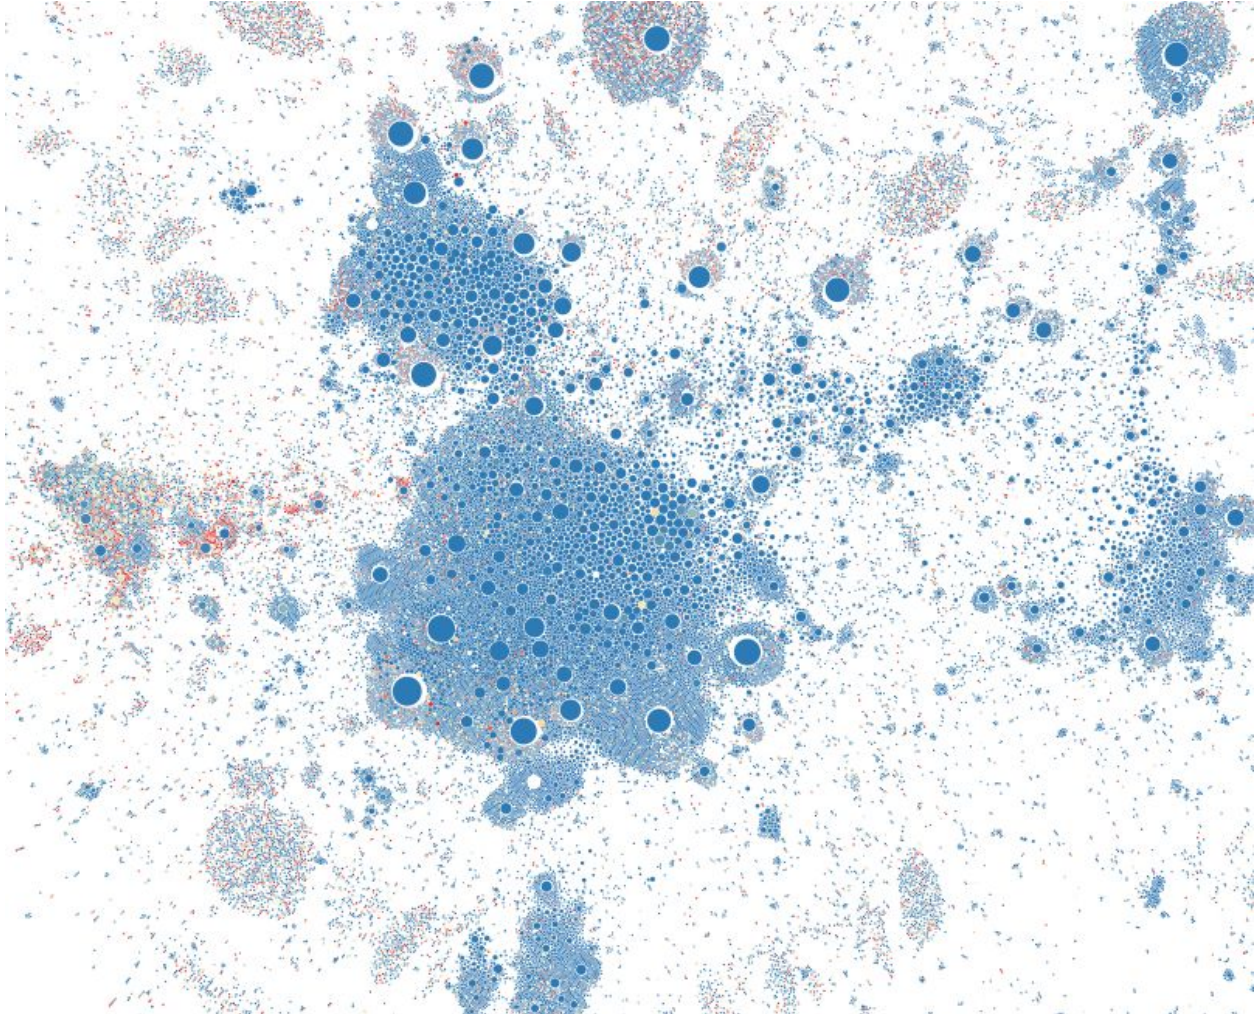
Figure 6: Detail of follow network with automation probabilities represented by colour from blue (0) to red (1).
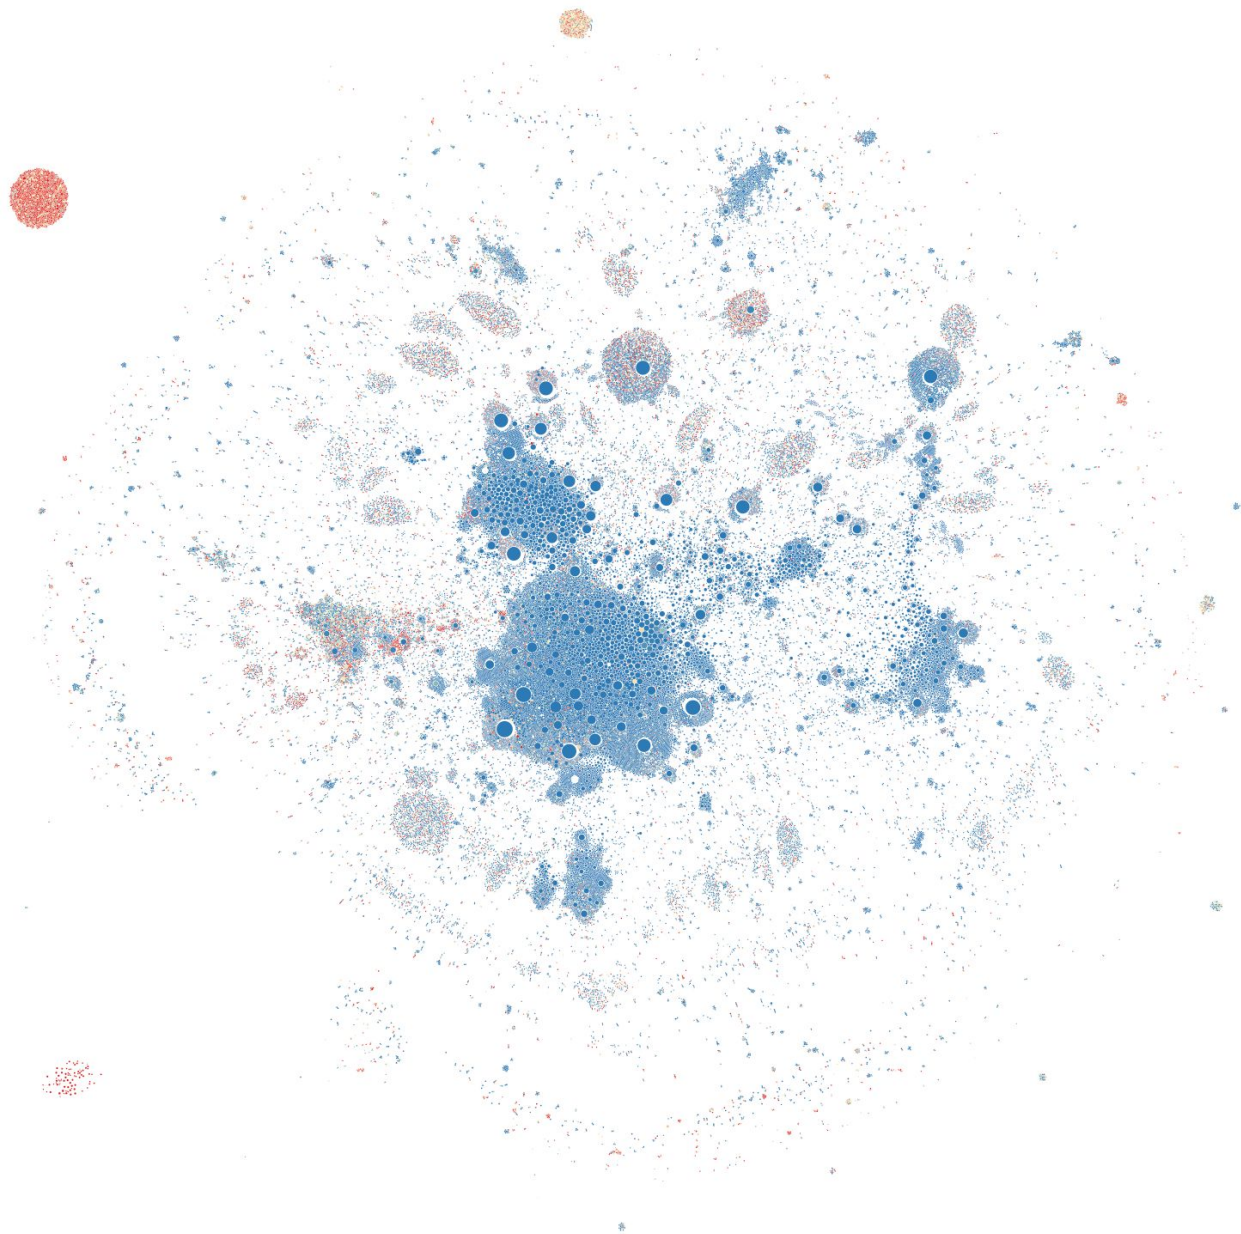
Figure 7: Overview of follow network with automation probabilities represented by colour from blue (0) to red (1). The red clusters in the upper and lower left corners have been identified as bot-nets.

References

Bruns, A., Moon, B., Münch, F. V., Wikström, P., Stieglitz, S., Brachten, F., & Ross, B. (2018). Detecting Twitter Bots That Share SoundCloud Tracks. In *Proceedings of the 9th International Conference on Social Media and Society - SMSociety '18*. Copenhagen, Denmark: ACM Press. http://doi.org/10.1145/3217804.3217923

Münch, F. V. (2019). *Measuring the Networked Public – Exploring Network Science Methods for Large Scale Online Media Studies*. Queensland University of Technology. Retrieved from https://eprints.qut.edu.au/125543

Salamanos, N., Voudigari, E., & Yannakoudakis, E. J. (2017). A graph exploration method for identifying influential spreaders in complex networks. *Applied Network Science*, *2*(1), 26. http://doi.org/10.1007/s41109-017-0047-y

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). http://doi.org/10.1038/s41467-018-06930-7

Varol, O., Ferrara, E., Davis, C. A., Menczer, F., & Flammini, A. (2017). Online Human-Bot Interactions: Detection, Estimation, and Characterization. In *International AAAI Conference on Web and Social Media*. Retrieved from https://aaai.org/ocs/index.php/ICWSM/ICWSM17/paper/view/15587/14817

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, *359*, 1146–1151. http://doi.org/10.1126/science.aap9559