Detecting Users Botness On Meetup.com

Samer Al-khateeb[0000−0001−6327−5720] and Cameron Kelly

Creighton University, 2500 California Plaza, Omaha, NE 68178, USA.
https://www.creighton.edu/
{sameral-khateeb1, cameronkelly}@creighton.edu

Abstract. Event-based social media sites (EBSMs) provide fertile ground for organizing and coordinating events, making them a valuable space for studying participants' characteristics, motivations, and roles in these events. One of the most popular EBSMs is Meetup.com, from which we collected data for this study. While many social media users are human, numerous studies have demonstrated the existence of automated programs that mimic human behavior on social media—a.k.a. social bots. In this research, we aim to identify accounts that exhibit such behavior on Meetup.com by leveraging various heuristics found in the literature and applying deep learning to assess the "botness" of these users. We also use Latent Dirichlet Allocation (LDA) topic modeling and ChatGPT to understand these bots' topics and roles on the platform. Results show that users can be classified by their "botness" with high accuracy, that factors beyond the heuristics significantly contribute to determining "botness", and that bots play various roles in event coordination and management.

Keywords: Social Bots · Meetup · Event-Based Social Media · Coordinated Events · Deep-learning · GPT · ChatGPT · LDA.

1 Introduction

With the widespread use of social media by "Gen Z" and "Gen Alpha" [24], organizing and coordinating events using this technology is becoming prevalent. Events are often coordinated using social media due to their large audience, ease of use, affordability, etc. Event-based social media sites (EBSMs) are specifically designed to facilitate such efforts. As of January 1, 2025, one of the most well-known EBSMs is Meetup.com. It allows a group of like-minded people to join "groups" online for an "event" that may occur in cyberspace, physical space, or both. Coordinated events involve planning and performing various tasks or steps to ensure their success in achieving a specific goal. A well-documented example of coordinated events is the formation of a "mob," which entails an underlying organizational process [3]. A mob is an event in which a group of people use social media to organize, coordinate, and even evaluate the impact of their event/mob. In a mob, people gather either online or offline to achieve a specific goal and then disperse. Mobs can be deviant (harmful and illegal), e.g., a mob attack on a Nordstrom store in California [8] and a 7-Eleven gas station

[7] or benign (fun and/or helpful), e.g., a "flash mob" of a group of people dancing in a shopping mall to entertain shoppers or to raise funds for a cause or charity [3]. Although most events—especially mobs—are attended by people (humans), social bots play a major role in organizing and coordinating these events. Various studies have highlighted the role of such actors on Twitter during flash mobs [4], natural disasters [19][20], large-scale combat operations [15], and ISIS propaganda-disseminating campaigns [2]. Other research highlighted the evolution of bots in disseminating propaganda during various military exercises [1] and compared the use of social bots on Twitter and Facebook [25]. However, to the best of our knowledge, no study has examined the role of bots in event-based social media (EBSM) platforms. Therefore, in this research, we collect data from Meetup.com and use deep learning techniques, topic modeling techniques, and large language model (LLM) to answer the following research questions:

1. What are some of the characteristics of social bots found in the literature?
2. What characteristics do Meetup bots exhibit? Furthermore, how can we rank the importance of these characteristics in differentiating bot accounts?
3. What *topics* and *roles* do bots play in Meetup events?

The rest of this paper is organized as follows: Section 2 provides a brief review of the literature related to this research. Section 3 details the methods used to collect and enhance our data, the deep learning model employed to classify Meetup users, as well as the topic modeling and LLM used to answer research question 3. Section 4 highlights our findings, while Section 5 concludes the study and suggests possible directions for future research.

2 Literature Review

We review recent research related to our work, focusing on articles published within the last 10 years, which we obtained from Google Scholar using the keyword "social bot". Literature review articles were given priority, as they provide a more comprehensive analysis of existing studies. While this is not an exhaustive approach, it offers a solid overview of relevant research. Notably, a search on Google Scholar using the keyword "*meetup bot*" yielded no articles among the top 10 pages (or top 100 results) that study or characterize Meetup bots, adding more novelty to this research. Due to paper length restrictions, we highlight only major findings here. Many researchers point out the ever-changing nature of social bots and how challenging they are to detect [13][18]. Several emphasize the need for automated methods rather than relying on human detection. While machine learning (ML) methods generally outperform human inspectors [12][26][27], deep learning appears to outperform all other approaches [14][16][21][23]. Many researchers have also attempted to develop automated tools and methods for detecting social bots on various social media platforms [6][9][13][17][26]. In this research, we leverage all these authors' findings and follow their suggestion to use deep learning models to detect user botness in one of the most used event-based social media sites, i.e., Meetup.com. A summary of the mentioned bot account characteristics/heuristics/traits is shown in Table 1.

3 Methodology

Below, we detail the method we used to conduct this research, starting with our data collection and enhancement techniques, followed by an explanation of the deep learning model used to classify the users and the SHAP library to rank the attributes importance in classification, and concluding with our use of Latent Dirichlet Allocation (LDA) and ChatGPT with GPT-4o to understand the topics and roles of the bot-ish users.

3.1 Data Collection & Enrichment

In this research, we leverage the Meetup.com data collected by Al-khateeb et al. [5]; however, in this research, we mainly use the data in the "users", "comments", and "replies" tables which have $9,864$ unique users who wrote $5,036$ comments and $2,267$ replies on $3,653$ events. For each user, we have the following data attributes: *userId, userName, bio, city, country, numEventsOrganized, numGroupsOrganized, uniqueness, numComments, numReplies, uniqueTopicsHosted, avgInteractionRate, numEventsAttendedWithMultiTopics, utility, timeTillEvent*. As a data enrichment step, we calculated 6 more data attributes, these are the criteria (see Table 1) that we found in the literature and we were able to estimate using our data to give "botness" score for each user. Each user account was assigned a score between 0 and 5, *zero* means the user did not have any of the criteria, while 5 means the user had all the criteria. Figure 1 shows the distribution of the accounts per criteria. It is worth mentioning that despite having a total of 6 criteria, there were no users that had all 6.

3.2 Deep-learning Model

We used the *Sequential()* model from the Keras API, integrated into the TensorFlow framework, to build and train our neural network. The network consists of an input layer with 128 neurons, two hidden layers with 64 and 32 neurons respectively, and a single output layer with 6 neurons. We divided our data into an 80% training-validation set and a 20% testing set. The training-validation set was further split into 70% for training and 10% for validation. We experimented with various parameter values; however, the best performance was achieved by training the neural network for 200 epochs using a batch size of 16 and 10-fold cross-validation. To improve training performance, stability, and speed, we incorporated batch normalization and used the LeakyReLU() activation function. To prevent overfitting, a dropout rate of 0.5 was applied after each layer, and early stopping was implemented. The model was trained on our lab computer which is equipped with a 2GHz 16-core Intel Xeon W processor (Turbo Boost up to 4.4GHz) and a Radeon Pro W5500X graphics card with 8GB of GDDR6 memory. Finally, we used the SHAP Python library [22] to perform an attribute importance analysis. The SHAP library can explain the outputs of machine learning and deep learning models by ranking how each feature contributes to classifying a sample into a specific class.

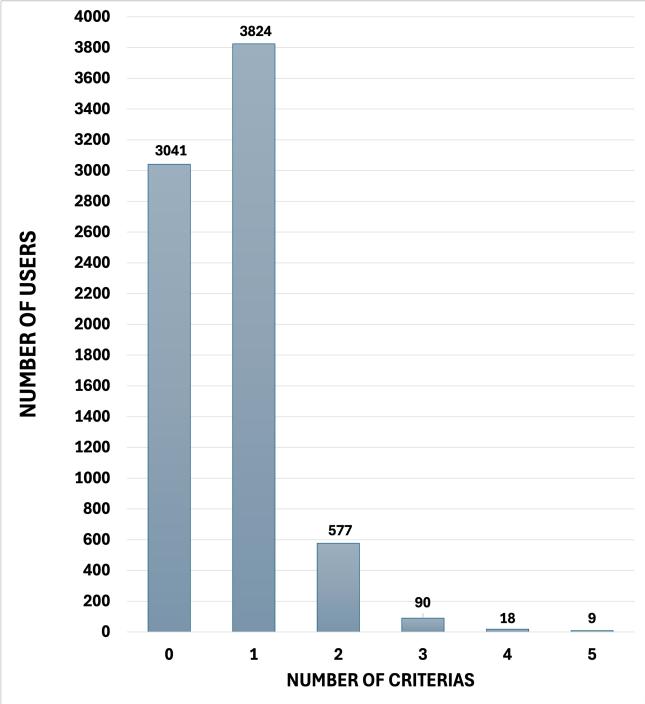

Fig. 1. Distribution of user accounts per criteria matching.

3.3 LDA Topic Modeling and ChatGPT with GPT-4o

We used Latent Dirichlet Allocation (LDA) [10], one of the most popular unsupervised topic modeling techniques, to identify *topics*—sets of words that best represent the themes of the provided documents—in the comments and replies of the users with the highest botness scores (i.e., users with botness scores of 4 and 5). We trained the model using a range of topic numbers and monitored the *coherence score*, which is used to evaluate the quality of the topics generated by the model (i.e., how interpretable or semantically consistent the topics are), to determine the optimal number of topics (i.e., the number of topics with highest coherence score). We then trained the model using this number and extracted the top five posts per topic. These posts best represent each topic, so we can manually examine them to identify the underlying themes of these bot-like accounts' comments and replies.

To further infer the possible *roles* these bots were playing in the events, we prompted GPT-4o as followed by [11] with the following: "*If these* _____ *were made by bots, what roles would the bots be playing?*" We filled in the blank with the word "comments" first, then in the second prompt with the word "replies" [11]. We analyzed the results, and the findings are reported below.

4 Findings

In this section, we report our findings, grouping them according to the research questions.

For the *first* research question, we reviewed multiple research articles, as outlined in Section 2, and identified approximately 21 criteria that have been recently used by various bot detection tools or highlighted by researchers as indicators of bot-like user behavior. These criteria are presented in Table 1.

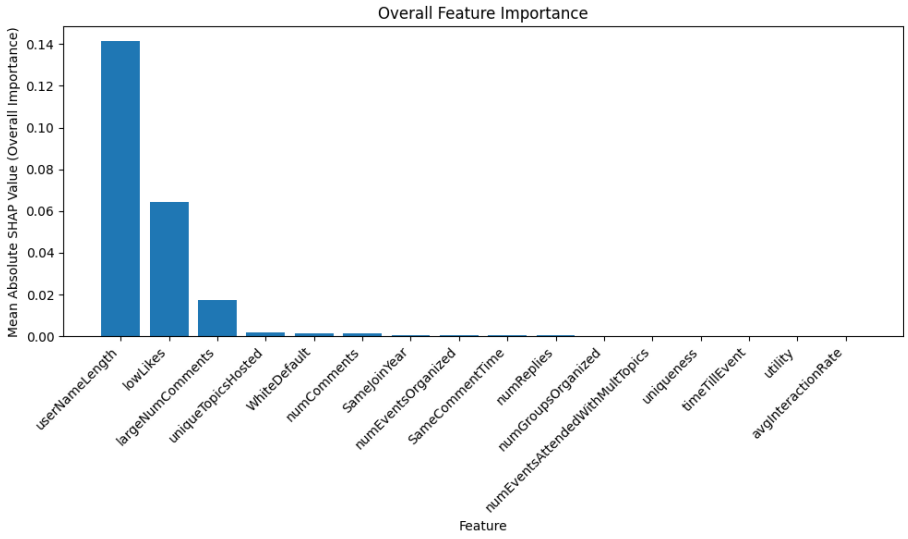

Fig. 2. Overall features importance.

For the *second* research question, we found that 6 (marked with an asterisk in Table 1) of the 21 identified criteria could be applied to the dataset we collected from Meetup.com. So, we estimated these criteria using various SQL queries and other techniques, then manually assigned each user a "botness" score ranging from 0 to 5 based on how many of these criteria the user met. This botness score also served as the six classes used to train our classifier to predict user botness. Our classifier achieved a high validation accuracy of 99.40% and a testing accuracy of 99.60%. Since the testing accuracy is slightly higher than the validation accuracy, our model is not overfitting. The high accuracy indicates that the users across the six classes have very distinguishable features that the model learned and can distinguish in the unseen testing data.

To determine feature importance, we used the SHAP library as explained in 3.2, and found that, overall, the top three features were *username length*, *low number of likes*, and *high number of comments* (see Fig. 2). For each individual

Fig. 3. Topics Per Comments.

Fig. 4. Topics Per Replies.

class (0–5), we found that our criteria were ranked highest in terms of importance, which makes sense since the botness score corresponds to the number of criteria matched. However, we also identified other important attributes that ranked among the top 10, aside from our criteria. These include the *number of comments*, *unique topics hosted*, *uniqueness*, *number of events organized*, *number of replies*, *number of events attended with multiple topics*, and *number of groups organized*.

For the *third* research question, we used LDA and ChatGPT with GPT-4o (as of April 7, 2025), as described in Section 3.3, to understand the *topics* these bots are engaging in and the possible *roles* they are playing.

LDA identified 7 and 5 as the optimal number of topics for the *comments* and *replies*, respectively. By visually inspecting the Intertopic Distance Map of both *comments* and *replies*, we found that the majority of the topics are well separated—except for Topics 2 and 5, and Topics 4 and 6 in the *comments* (see Fig. 3), and Topics 4 and 5 in the *replies* (see Fig. 4)—indicating similarity between those topics. Note that the size of the circles in both figures reflects how common each topic is across the corpus. For the *comments* topics, we found the following: Topic 1 is about *preparation and participation updates*, Topics 2 & 5 are *communication with the organizer*, Topic 3 is about the *readiness of participants*, Topics 4 & 6 are *information-seeking and participation updates*, Topic 7 is *event logistics, such as transportation updates and personal engagement*. For the *replies* topics, we found: Topic 1 is about *group participation encouragement*, Topic 2 is *information on event planning and activities*, Topic 3 is *personal updates*, Topic 4 shows *participants excitement and motivation*, and Topic 5 has *additional personal updates*, possibly overlapping with Topic 3.

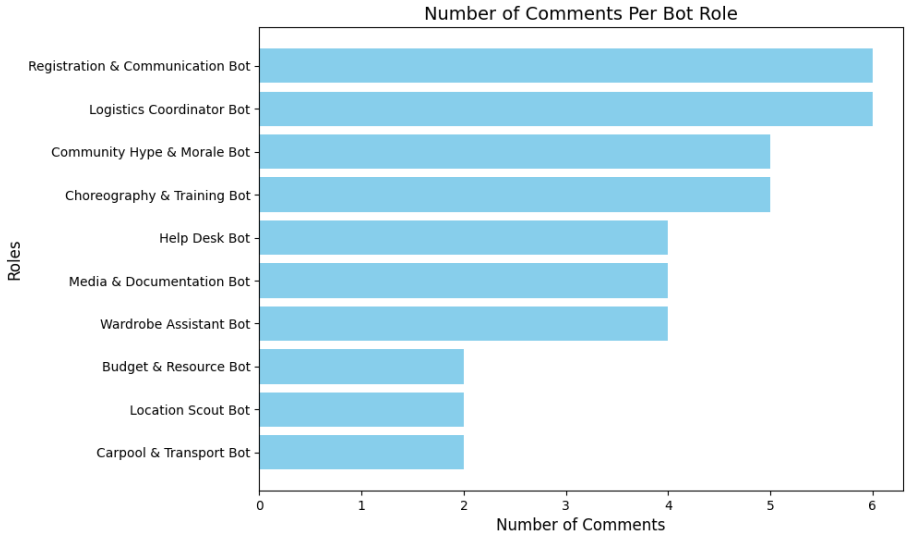

Fig. 5. Top Comments Per Bot Role.

Using ChatGPT and the top five *comments* from each topic identified by LDA, we found that, in general, these bots appear to play various *support roles in event coordination or management*. To better understand these roles, we prompted ChatGPT to break them down into sub-roles based on the provided comments. The results revealed that most *comments* focused on event and training logistics, such as wardrobe, budget, resources, location, and transportation (see Fig. 5). We also analyzed the top five *replies* from each topic identified by LDA to further understand the bots' roles. These *replies* typically *confirmed attendance, posted motivational messages, or asked the event organizers questions* (see Fig. 6). Notably, none of these comments or replies appeared "bot-ish," to us even though they were posted by accounts with high botness scores—that is, accounts that met all (5) or nearly all (4) of our user botness criteria. This may suggest either a very high level of bot quality (i.e., bots that are difficult to distinguish from humans) or that the criteria we used may not be sufficient to reliably identify bots. We are leaning toward the first scenario, especially since the analysis here used text (comments and replies), which can be easily generated by LLMs. However, further investigation is needed to explore both possibilities.

5 Conclusion & Future Research

In this research, we examined how "bot-ish" users behave on one of the most popular event-based social media platforms, Meetup.com—an understudied problem in the literature. We identified bot traits from existing research and manually

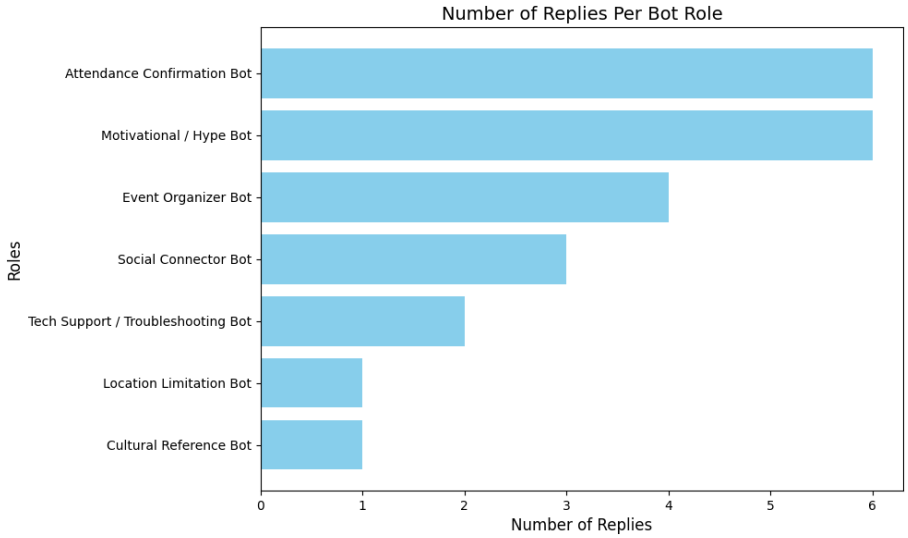

Fig. 6. Top Replies Per Bot Role.

scored users in our dataset from 0 to 5, based on the number of criteria they met. Then, we trained a deep neural network to recognize patterns for each class and used it to predict user botness. The classifier achieved very high accuracy on the unseen test data. Results also show that three of the six criteria we estimated are among the most important features for determining user class. We further investigated the role of these bot-ish users in coordinated events using LDA topic modeling and ChatGPT. The results suggest that bot-ish users engage in legitimate event coordination activities rather than typical malicious behaviors, and these users play various roles in *event coordination and management*. Further analysis is required to determine whether these users exhibit strong bot-like behavior and are, in fact, bots, or if they are human. One possible direction for future research is to use unsupervised ML to cluster users based on their similarities (without using predefined criteria) and compare the resulting clusters with our assigned classes.

Acknowledgments. This work is based upon work supported in part by the Office of the Under Secretary of Defense for Research and Engineering (FA9550-22-1-0332) and the U.S. Army Research Laboratory (W911NF-25-1-0147). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. Department of Defense.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article. The authors declare the funding received to conduct this research in the acknowledgments section.

References

1. Agarwal, N., Al-Khateeb, S., Galeano, R., Goolsby, R.: Examining the use of botnets and their evolution in propaganda dissemination. Defence Strategic Communications **2**(1), 87–112 (2017)
2. Al-Khateeb, S., Agarwal, N.: Examining botnet behaviors for propaganda dissemination: A case study of isil's beheading videos-based propaganda. In: 2015 ieee international conference on data mining workshop (icdmw). pp. 51–57. IEEE (2015)
3. Al-khateeb, S., Agarwal, N.: Flash mob: a multidisciplinary review. Social Network Analysis and Mining **11**(1), 97 (2021)
4. Al-khateeb, S., Anderson, M., Agarwal, N.: Studying the role of social bots during cyber flash mobs. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation. pp. 164–173. Springer (2021)
5. Al-khateeb, S., Burright, J., Fernandes, S.L., Agarwal, N.: Analyzing and predicting meetup mobs outcome via statistical analysis and deep learning. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation. pp. 164–173. Springer (2024)
6. Aljabri, M., Zagrouba, R., Shaahid, A., Alnasser, F., Saleh, A., Alomari, D.M.: Machine learning-based social media bot detection: a comprehensive literature review. Social Network Analysis and Mining **13**(1), 20 (2023)
7. Arreola, A., Lloyd, J.: Watch: Street takeover 'flash mob' swarms los angeles 7-eleven (2022), https://www.nbclosangeles.com/news/local/watch-unruly-street-takeover-crowd-swarms-los-angeles-7-eleven/2967174/
8. Barnard, C.: Brazen flash mob-style robbery of walnut creek nordstrom sparks outrage from city leaders (2021), https://abc7news.com/walnut-creek-nordstrom-looting-robbery-smash-and-grab-wc-broadway-plaza/11260247/, section: crime-safety
9. Ben, N.: #BotSpot: Twelve Ways to Spot a Bot (Sep 2017), https://medium.com/dfrlab/botspot-twelve-ways-to-spot-a-bot-aedc7d9c110c
10. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. Journal of machine Learning research **3**(Jan), 993–1022 (2003)
11. Burright, J., Al-khateeb, S.: A comparative analysis of the ethics of gene editing: Chatgpt vs. bard: J. burright, s. al-khateeb. Computational and Mathematical Organization Theory **31**(2), 195–206 (2025)
12. Cai, M., Luo, H., Meng, X., Cui, Y., Wang, W.: Network distribution and sentiment interaction: Information diffusion mechanisms between social bots and human users on social media. Information Processing & Management **60**(2), 103197 (2023)
13. Cresci, S.: A decade of social bot detection. Communications of the ACM **63**(10), 72–83 (2020)
14. Ferrara, E.: Social bot detection in the age of chatgpt: Challenges and opportunities. First Monday (2023)
15. Galeano, R., Galeano, K., Al-Khateeb, S., Agarwal, N., Turner, J.: Botnet evolution during modern day large-scale combat operations. Perceptions Are Reality: Historical Case Studies of Information Operations in Large-Scale Combat Operations pp. 163–173 (2018)
16. Hayawi, K., Saha, S., Masud, M.M., Mathew, S.S., Kaosar, M.: Social media bot detection with deep learning methods: a systematic review. Neural Computing and Applications **35**(12), 8903–8918 (2023)

17. IONOS, e.t.: Social bots – the technology behind fake news (Mar 2022), https://www.ionos.com/digitalguide/online-marketing/social-media/social-bots/
18. Kenny, R., Fischhoff, B., Davis, A., Carley, K.M., Canfield, C.: Duped by bots: why some are better than others at detecting fake social media personas. Human factors **66**(1), 88–102 (2024)
19. Khaund, T., Al-Khateeb, S., Tokdemir, S., Agarwal, N.: Analyzing social bots and their coordination during natural disasters. In: Social, Cultural, and Behavioral Modeling: 11th International Conference, SBP-BRiMS 2018, Washington, DC, USA, July 10-13, 2018, Proceedings 11. pp. 207–212. Springer (2018)
20. Khaund, T., Bandeli, K.K., Hussain, M.N., Obadimu, A., Al-Khateeb, S., Agarwal, N.: Analyzing social and communication network structures of social bots and humans. In: 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM). pp. 794–797. IEEE (2018)
21. Li, S., Yang, J., Zhao, K.: Are you in a masquerade? exploring the behavior and impact of large language model driven social bots in online social networks. arXiv preprint arXiv:2307.10337 (2023)
22. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. Advances in neural information processing systems **30** (2017)
23. Lyu, N., Xu, B., Guo, F., Shen, H.: Dcgnn: Dual-channel graph neural network for social bot detection. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. pp. 4155–4159 (2023)
24. Najib, S.: Gen Beta kicks off in 2025: Your guide to all the generation names and years (Dec 2024), https://abcnews.go.com/GMA/Living/generation-names-and-years/story?id=114802892
25. Obadimu, A., Mead, E., Al-Khateeb, S., Agarwal, N.: A comparative analysis of facebook and twitter bots. In: SAIS 2019 Proceedings. pp. 1–6. Southern Association for Information Systems (SAIS) (2019)
26. Wu, J., Ye, X., Mou, C.: Botshape: A novel social bots detection approach via behavioral patterns. arXiv preprint arXiv:2303.10214 (2023)
27. Zhang, Y., Song, W., Shao, J., Abbas, M., Zhang, J., Koura, Y.H., Su, Y.: Social bots' role in the covid-19 pandemic discussion on twitter. International Journal of Environmental Research and Public Health **20**(4), 3284 (2023)

Table 1. Heuristics from the literature are listed below; an asterisk (*) by the Heuristic No. indicates it was estimated and used in our deep learning model.

Num	Heuristic Explanation	Platform Used	Reference
1	Use of URL shortener	Twitter	[6][9]
2	High following to follower ratio	Twitter	[6][14][17]
3∗	Young account age and many posts (we used same join year)	Twitter	[6][26]
4	Delay in posting on current events	Twitter	[21][27]
5∗	Low number of likes on a post (less than average)	Twitter, Facebook	[9][6][17][27]
6	Post repetitive content	Twitter	[6][14]
7∗	Large number of comments/post	Twitter	[9][26]
8	50 posts per day is suspicious and 144 posts per day is highly suspicious	Twitter	[9]
9	Less personally identifying information	Twitter	[9]
10∗	Stolen or shared profile picture (we used white or default profile pictures)	Twitter	[9][13]
11	Certain word collocation in account content such as "news", "world", "com", "account", "community", "people"	Twitter	[27]
12	Users who have more re-posts than posts are more likely to be bots	Twitter	[12]
13	Bots lag when posting about emerging events	Twitter	[21]
14∗	Length of screen name	Twitter	[26]
15∗	High frequency of activities and messages (we used multiple comments at the exact same time)	Twitter	[14]
16	High usage of hashtags	Instagram	[6]
17	Disabled geographic location	Twitter	[6]
18	Less diversity in geographic locations	Twitter	[27]
19	More neutral or positive sentiments regarding public events compared to humans	Weibo, Twitter	[12]
20	Often unverified accounts	Twitter	[6]
21	The same user repeatedly uses the same set of keywords in many of their posts	Twitter	[6][14][21]