

A Discussion of Detection of Mutual Influences between Socialbots in Online (Social) Networks

Stefanie Urchs
University of Passau
Passau, Germany
Stefanie.Urchs@uni-passau.de

Abstract—Many people organise themselves online in social networks or share knowledge in open encyclopaedias. However, these networks do not only belong to humans. A huge variety of socialbots that imitate humans inhabit these and are connected to each other. The connections between socialbots lead to mutual influences between them. If the influence socialbots have on each other are too big they adapt the behaviour of the other socialbot and get worse in imitating humans. Therefore, it is necessary to detect when socialbots are mutually influencing each other. For a better overview socialbots in the social networks Facebook, Twitter and in the open encyclopaedia Wikipedia are observed and the mutual influences between them detected. Furthermore, this paper discusses how socialbots could handle the detected influences.

Index Terms—organic computing, social networks, mutual influences, social bots

I. INTRODUCTION

In the year 2018 the number of social media users increase by 13 percent to 3.2 billion users¹. More and more people connect with each other over these networks, share knowledge or promote themselves. Even companies have discovered online networks for marketing and politicians for campaigning for their cause [1]. With the increased usage of these networks humans handed tasks over to socialbots. These socialbots now campaign for companies and politicians or edit articles on online encyclopaedias like Wikipedia. To do so the socialbots often connect with other users. But not all of these users are human, therefore, do socialbots connect to socialbots and try to influence them. Furthermore, socialbots adapt their behaviour according to feedback they get from other users in the online (social) networks. However, some of these users are socialbots. If a socialbot adapts too much to the feedback of other socialbots it becomes less and less human like. A socialbot that does not act like a human is easily detected as a socialbot and either deleted by the online network it inhabits or avoided by other users.

This paper defines socialbots as autonomous agents, that imitate human behaviour (section II-C). Socialbots are, therefore, self-organising and self-adapting systems. These two terms belong to the field of Organic Computing (section II-A). The term self-organisation describes the process of

satisfying objectives without or with minimum external intervention [2]. In turn the term self-adapting, also called self-configuration, describes the modifications the system makes on its own parameters to reach higher-level user goals [3].

In this paper socialbots are considered as Organic Computing systems. Therefore, a taxonomy for detecting mutual influences in organic computing systems [4] is used to explore mutual influences between bots.

The environment in which a socialbot exists defines how it could influence other socialbots. Therefore, the main contribution of this paper is to use the above taxonomy in different online (social) networks. Further, it is discussed how bots could handle mutual influences. For this paper the two popular social networks Facebook² and Twitter³ are chosen. Additionally, the open encyclopedia Wikipedia⁴ is examined. Facebook has worldwide the most active users⁵. On Facebook a user owns an account and creates a profile about themselves. On this profile they can share their name, gender, interests and contact information. User form links to other users. This links represent social relationships like friendships or work acquaintances [5]. Moreover, Facebook is used for opinion distribution. For example, politicians use Facebook to promote their election campaigns [6].

Twitter is a messaging service where people can share short messages anonymously with the world. The usage of bots on twitter is relatively easy: to register an account one only needs to provide an email-address and a mobile phone number, followed by passing a CAPTCHA recognition. In addition, the Twitter API makes it possible to automate actions on Twitter [7].

On Wikipedia users can share their knowledge. The content is moderated by a community of users. The ability to fully moderate articles is linked to an registered user account⁶. To help the human moderators it is possible to create bots. These bots have to be registered on Wikipedia and are flagged as

²<https://www.facebook.com/> (accessed on 15.12.18)

³<https://www.twitter.com/> (accessed on 15.12.18)

⁴<https://www.wikipedia.org/> (accessed on 15.12.18)

⁵<https://www.statista.com/statistics/272014/>

global-social-networks-ranked-by-number-of-users/ (accessed on 15.12.18)

⁶https://www.en.wikipedia.org/wiki/Wikipedia:Why_create_an_account (accessed on 15.12.18)

¹<https://www.smartinsights.com/social-media-marketing/social-media-strategy/new-global-social-media-research/> (accessed on 15.12.18)

bots and officially approved [5].

This paper answers the following research questions:

How do socialbots, in online networks, influence each other?

How can these influences be detected and handled?

The remaining paper is structured as follows: the second section introduces the related research in the fields of Organic Computing, Online Social Networks and Socialbots. Subsequently, the third section describes mutual influences in Organic Computing systems. The idea of mutual influences is continued in the fourth section where mutual influences of socialbots in Facebook, Twitter and Wikipedia are discussed. The fifth section reviews how bots could handle the mutual influences. In the sixth and last section the paper is concluded.

II. RELATED RESEARCH

This paper relies on research that is already done in the fields of Organic Computing, Online Social Networks and Online Social Bots. Therefore, this research is discussed in this chapter.

A. Organic Computing

In recent years the complexity of computer systems increased quite fast. In order to handle this increasing complexity Tomforde et al. [3] describe the term of Organic Computing (OC). They define OC systems as technical systems that perceive their environment with hardware and virtual sensors and manipulate their environment through actuators. These systems are able to dynamically and autonomously adapt to the environment. Every time a system adapts to the environment its utility is affected, therefore, is the system able to improve its utility continuously. This adaptive behaviour is possible through the utilization of so-called self-* mechanisms. A typical OC system consists of various autonomous organic subsystems. To handle these subsystems consist OC systems out of two complimentary parts: the first one operates the system and the second one is responsible for the adaptation processes. OC systems are normally based on machine learning techniques, to be able to react appropriate to unknown and unanticipated conditions. Organic Computing aims to enhance technical systems with properties that can be found in alive things. This does not mean to build systems out of organic tissue but to transfer the behaviour found in nature to technical systems.

The term OC occurs first in 2004 in an article of Christian Müller-Scholer et al. [8] about OC. In more recent publications Tomforde joined Müller-Scholer, therefore, mostly publications of Tomforde et al. are used.

B. Online (Social) Networks

In this paper online networks are defined as online places where people interact in any kind of way. The network between the users is formed by the interaction between the users.

Online Social Networks are platforms that are hosted online. To fully participate in such a platform users register accounts

on these platforms. The registration is done under the real name or a pseudonym. Users can enter personal information about themselves which is added to their profile. To form a network users can connect their accounts to the accounts of other users. The connection between these accounts are formed because of various reasons: real-world/online friendships, work relations, shared interests or interest in the contribution of the other person. Furthermore, it is mostly possible for users to join or create groups where they can share messages and content [9].

C. Socialbots

Bots are autonomous software agents that act in place of a human. Franklin and Graesser define autonomous agents as follows "An autonomous agent is a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future." [10]

M. Tsvetkova et al. [5] define online bots as bots run in the internet. Because of a huge variety of different kinds of bots online they categorise online bots in the following four categories:

- collect information
- execute actions
- generate content
- emulate humans

In all categories they detect benevolent and malevolent bots. However, they argue that bots are not capable of "emotions, meaning-making, creativity and sociality" [5]. Therefore, the bots themselves are not bad or good but the way they are used is.

In the context of online social media so called socialbots can be found. According to Boshmaf et al. [11] can socialbots be distinguished from other bots by the attempt of socialbots to seem human.

III. MUTUAL INFLUENCE

This section discusses how mutual influences can be measured and detected in OC systems.

To define mutual influences a model of the system in which these influences occur has to be defined. Rudolph et al. [12] define such a system model as follows: the system consist of a set of agents where each agent can take on different configurations. Each configuration is comprised of different parts. For example, a router can consider different configurations like the processed network protocol or parameter settings. Moreover, it can be assumed that the configurations of the agents are not overlapping. However, this does not imply that all configuration parts have to be disjoint. For example, may two routers be able to configure the time out. This leads to the same set of configurations on two different devices. Further, a local performance measurement is needed to apply mutual influence detection. Each agent needs to validate the success of its actions at runtime. For this validation either feedback from the environment or manual

assignments are used.

Based on this system Rudolph et al. [12] introduce a measurement to gauge mutual influences. It uses stochastic dependency measures which estimate connections between the performance of an agent A and the configuration parts of an agent B. In order to do so the performance of A and B are each seen as random variables. The dependency measures are then used to find correlations between these random variables. If the correlation between the random variables is high the influence between these two agents is also high. This is because, the configuration of one agent matter for the performance of the other. If the correlation between the random variables is low the influence is also.

This paper uses the "taxonomy for Organic Computing systems regarding mutual influences" by Rudolph and Tomforde [4]. They introduce three important characteristics of OC systems: *entities*, *communication* and *influence*.

Entities

OC systems can be categorised based on the number of entities they contain. Small systems contain few entities, middle-size systems contain up to a few hundred entities and everything above is considered as large-scale systems. The type and the number of the configuration parts is also interesting. There are three types of configuration parts: nominal, ordinal or infinite real-valued. For nominal and ordinal configuration types the number of categories can be classified [4].

Communication

For the taxonomy two communication boarder cases are considered. First all agents utilise for example the same hardware and therefore the communication between the agents is free. Second agents are only able to communicate with their neighbours. Between these boarder cases exists various other communication possibilities [4].

Influence

Influence can manifest in the context of communication: when high communication costs exists more central agents have more influence than agents that can only be reached over several steps. A different way to observe influence is in the interaction of different agents. For example imagine a drilling robot with two robotic arms that hold a piece of wood while a third arm drills a hole in the wood. If each arm is regarded separately no influence can be revealed. The strength of the influence is measured with dependency measures. They can either be linear, monotonic or stochastic. Furthermore, can influence be noticeable either instantly or with a delay [4].

IV. MUTUAL INFLUENCE BETWEEN SOCIALBOTS

Mutual Influences are introduced generally above. This section describes how these can be detected in Facebook, Twitter and Wikipedia. These networks are chosen because, Facebook has

worldwide the most active users⁷. The Twitter API encourages the use of bots on Twitter. On Wikipedia everyone can alter articles, which has potential for socialbot usage. It is generally necessary to detect the influences between socialbots to prevent them from adapting to other socialbots. In the worst case a chain reaction starts, where socialbots influence each other again and again and the human input is totally discarded. To be able to detect other socialbots they need a concept of socialbots. For this concept certain characteristics of socialbots have to be defined. The source of every incoming influence has to be verified, regarding to the characteristics of socialbots. If the influencing entity is recognised as socialbot the received feedback should be ignored.

A. Facebook

Most research regrading bots on Facebook reviews the influence that content, generated by bots, has on world politics [13]. A different research field for bots on Facebook is the Facebook Messenger⁸. Relevant for this paper is the research about infiltrating Facebook with a socialbot network. Boshmaf et al. [11] did this in 2013. Their socialbot network consists of three components: several socialbots that each own a profile on Facebook, a botmaster and a command and control channel. Socialbots can either interact socially (for example posting a message to their time-line) or in a way concerning the structure of the network (for example sending a friend request to another user). The interactions a socialbot does are either predefined locally or send by the botmaster. Further, socialbots gather data about users (so called botcargo) and send it to the botmaster. The botmaster is a controller that can be accessed by a human. All communication between socialbots, botmaster and human happen via the command and control channel.

The proposed socialbot network is centrally controlled and the socialbots do not act autonomously. In addition, no local performance measure is proposed. Therefore, the above socialbot network is adjusted as follows: the new bot network follows an observer/controller framework [14]. Each socialbot serves as observer and gathers data from the environment. With this data a situation description is build and send to a controller. Moreover, a socialbot should be able to post to its wall, send friend request (to get connections to neighbouring entities), accept/deny friend requests, end friendships and like pots of users it is "friends with" (a socialbot is friend with users that accepted its friend request). To keep the notation of Boshmaf et al. this controller is called botmaster. The botmaster selects an action based on the situation description and sends this action back to the socialbots. Furthermore, the botmaster evaluates the success of the last action depending on the reaction of the environment. This success could be measured in likes a post gets, in friend request that get accepted and how many/long friendships can be maintained. Note that all socialbots got an individual botmaster, they are

⁷<https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/> (accessed on 15.12.18)

⁸<https://www.facebook.com/messenger/> (accessed on 04.03.19)

not controlled centrally any more.

1) *Entities*: The bot network Boshmaf et al. [11] created consisted of 102 socialbots. Therefore, it can be classified as a middle-sized system. Possible configuration for a socialbot are in which manner it is posting to its wall, to which users it sends friend requests (preferentially to users who already share friends or not) in which intervals/frequency it posts to its wall, when friend request are accepted and when friendships are terminated. Most types of these configurations are nominal. For example, the manner of posting could describe the political direction, pro/anti religion, pro/anti animals, etcetera. The interval/frequency the socialbot posts to its wall is infinitely real-valued.

2) *Communication*: All socialbots live in the same system, therefore, the communication costs are too small to be relevant and can be neglected.

3) *Influence*: The influence of a socialbot is not restricted to only a neighbour. On the contrary, every socialbot can influence every other entity in the whole network. As soon as a socialbot interacts with other users, it influences them. This includes sending/denying/accepting friend requests and liking the posts of other users. Even the text a socialbot posts on its wall can influence other bots. This is the case when the generation of new text is actively learned from new text posted in the network of the socialbot.

The centrality of a socialbot can be neglected because of the infinitesimal communication costs.

It is unlikely that the influence between the socialbots can be detected by a linear or monotonic measure. Therefore, a stochastic dependency measure should be used.

Cases in which the influence is only revealed if several socialbots work in common could exist. However, it is assumed that the influence is mostly directly revealed.

The influence is only revealed as soon as a different entity reacts to the actions taken by the socialbot. Therefore, is the influence noticeable with a delay.

B. Twitter

Plenty of research about bots on Twitter exists. Many researchers deal with political influences of bots [15]. Other try to detect bots on Twitter. For the later the so called DARPA Twitter Bot Challenge [16] was hold in February/March 2015. Multiple teams were challenged to detect bots based on previously detected ones.

Freitas et al. [17] chose a different approach. They build 120 socialbots and observed them, to get a better understanding of infiltration strategies in Twitter.

The socialbots can follow other users (get to know neighbouring entities), tweet (post a message, that is automatically generated with a Markov chain), and retweet Tweets (post a message another user already posted again) of users the socialbots are following. The socialbots are active in predefined, random intervals and are categorised as

high active/low active. The category influences the size of the random time intervals between two actions. In addition, socialbots "sleep" a predefined time, to create the impression of a sleeping human.

These socialbots do not act autonomously and do not learn from their behaviour. Therefore, the following adjustment is proposed: Socialbots should follow a observer/controller framework [14]. The socialbot observes its environment by measuring how many user interact with its Tweets (liking them, favouriting them or retweeting them). Furthermore, how many user follow the socialbot and how often other user mention the socialbot in their own tweets is measured. Additionally, should socialbots be able to like and favourite a Tweet.

1) *Entities*: Freitas et al. [17] created 120 socialbots, therefore, it can be categorised as middle-sized.

Possible configurations for the entities are in which intervals to take action (post a Tweet or follow a user), which kind of content to post (political direction, pro/anti religion, pro/anti animals, et cetera), which kind of users to follow (randomly chosen users/users that follow the same users as the socialbot). All above mentioned configurations are nominal.

2) *Communication*: As in the Facebook example: all socialbots live in the same system. Therefore, it possible to neglect the communication costs.

3) *Influence*: Like in Facebook the influence of a socialbot is not restricted to its direct neighbours. All socialbots can influence all entities within Twitter. The influence manifests in interactions with other users and their Tweets. This includes liking, favouriteing and retweeting a Tweet and following other users. Even the creation of a new Tweet can influence other socialbots, if they take this text to learn to generate new text. The centrality of an entity can be ignored because of the infinitesimal communication costs.

The revelation of the influence is assumed to be mostly direct. However, it is possible that collaboration between socialbots also reveals influence. Though, these cases should be less common.

To measure the dependency between socialbots stochastic measures should be used. It seems unlikely that linear or monotonic measures can capture the dependencies.

The influence a different entity has on the socialbot is noticeable when the second entity reacts to an action of the social bot. Therefore, a delay between action of the socialbot and the influence occurs.

C. Wikipedia

Seemingly no one has done research in infiltrating Wikipedia with a network of socialbots. This might be explained by a

strict blocking⁹/banning¹⁰ policy of Wikipedia. Due to these policies a consensus of the community is enough to ban an account or a IP-address from editing. All user who edit an article without an account have to reveal their IP-address¹¹, but to create an account only an username and a password is necessary¹². Furthermore, Wikipedia states: "Shared IP addresses such as school and enterprise networks or proxy servers are frequently blocked for vandalism..."¹³. These frequent blocks increase the difficulty of an infiltration.

A theoretical network of socialbots should first acquire multiple IP-addresses outside of a research institute. This should prevent a block/ban of the whole network at once. Second, should every socialbot get an user account, so IP-addresses are not revealed to the public. By hiding the IP-address it is no longer possible to easily look up the owner of an IP-address. Thirdly, every socialbot should be able to only edit articles. Bots could get more rights after a four day period. However, these advanced activities lead too a higher visibility in the community and, therefore, to a higher risk of detection. Every socialbot should follow a observer/controller framework [14]. The environment is modelled by measuring how long a change in an article is accepted (not rolled back to the initial article) and how long the socialbot is able to operate (no ban/block). The controller decides the following parameters:

- Which randomly chosen article to edit.
- Which parts of the article to edit.
- In which time intervals to act.

The changes in the article can be created by using an arbitrary natural language processing model which is trained on the whole of Wikipedia. The longer a change made by a socialbot stays in the system and the longer the socialbot stays active the higher the reward. The social bot gets to know its neighbouring entities by interacting with them. Whenever the socialbot changes an article this article was written and possibly changed by different entities. Moreover, the changes a socialbot makes to an article are possibly changed by a different entity. Finally, are blocks and bans imposed by neighbouring entities.

1) *Entities*: It seems that shared IP-addresses are at a high risk of being blocked by Wikipedia, therefore, should a socialbot network consist of unique IP-addresses. This restricts the amount of entities to a small sized system. Possible configurations are which article to edit, where this article should be edited and in which time intervals the socialbot should act. The first two configurations are nominal

⁹https://en.wikipedia.org/wiki/Wikipedia:Blocking_policy (accessed on 06.01.19)

¹⁰https://en.wikipedia.org/wiki/Wikipedia:Banning_policy (accessed on 06.01.19)

¹¹https://www.en.wikipedia.org/wiki/Wikipedia:Why_create_an_account (accessed on 06.01.19)

¹²<https://en.wikipedia.org/w/index.php?title=Special:CreateAccount> (accessed on 06.01.19)

¹³https://www.en.wikipedia.org/wiki/Wikipedia:Why_create_an_account (accessed on 06.01.19)

and the last one is infinitely real-valued.

2) *Communication*: According to Facebook and Twitter: all entities exist in the same system. Therefore, it is possible to neglect the communication costs.

3) *Influence*: A socialbot is able to influence every socialbot in Wikipedia by changing an article another socialbot previously changed.

The collaboration between socialbots should not lead to a different influence, as viewing the influences on their own. Communication costs are negligible and so is the centrality of an entity. Linear and monotone dependency measures do not seem to capture the full dependency between socialbots, therefore, stochastic measures are proposed.

As soon as a different entity changes the what the socialbot edited, the socialbot is affected. But the influence does not occur as soon as the initial socialbot takes action. Therefore, the influence is only noticeable with a delay.

D. Discussion

All three examples above are part of the same domain. Therefore, are the results discovered with the taxonomy relativity similar. To better differentiate between the online networks the taxonomy needs adjustment.

It is generally possible to measure if socialbots and different entities are influencing each other. To limit this detection to socialbots, it is necessary that all influences are categorised whether they come from a fellow socialbot or not. As stated above, socialbots need a concept of what determines a fellow socialbot to be able to detect them. The influences of other socialbots are not desirable, therefore, feedback from other socialbots should be ignored.

V. HANDLING OF MUTUAL INFLUENCES

The detection of mutual influences between online bots is only the first part. In the following section the next part is engaged by discussing how bots could handle the mutual influences.

A. Facebook

The way a socialbot A posts on its wall can be influenced by other socialbots by the number of likes they give A. Many likes indicate that the current strategy should be continued. This influence can be handled by determining if a like possibly comes from a socialbot or another entity. If the source is categorised as socialbot the like should be ignored.

B. Twitter

The frequency in which a socialbot tweets can be influenced by the interaction other socialbots have with the Tweets of the initial socialbot. If many entities like/favourite or retweet a Tweet the socialbot can increase the tweet-frequency. Fewer likes indicate that the socialbot tweeted too much. The influence other socialbots have can be restricted by ignoring all interactions they make.

C. Wikipedia

The length of the text a socialbot is changing depends on feedback from others. If the change is reverted, changed in a different way or the socialbot is banned/blocked the length of changes should decrease. To avoid influence coming from other socialbots reverts, changes blocks and banns from possible other socialbots should be ignored.

D. Discussion

The handling of influences from other socialbots is the same in all instances. All received feedback should be categorised in feedback from socialbots and feedback from other entities. Feedback received from socialbots should be ignored to prevent the adaptation of behaviour stemming from other socialbots.

VI. CONCLUSION AND FUTURE WORK

This paper introduced a taxonomy for Organic Computing systems regarding mutual influences. This taxonomy was applied to socialbots in online networks. Furthermore, it was discussed how these influences could be handled.

The chosen taxonomy is generally able to detect influences in the chosen examples. However is it too general to differentiate between the different online networks.

Influences coming from other socialbots should generally be ignored to prevent the adaption to the behaviour of other socialbots. To be able to ignore other socialbots a socialbot needs criteria which define a socialbot. It then needs to categorise all feedback in coming from a socialbot or not.

In the future a taxonomy specific to online networks could be developed. Moreover, the criteria that define a socialbot should be determined independently for every online (social) network. Generally should the research of autonomous bot networks in online networks be promoted. Especially the research on socialbots in Wikipedia. A possible additional future research direction is to look into automatic stock trading bots. These systems are highly interwoven, because of the stock market they are inhabiting. Small changes of one system can lead to a catastrophic chain-reaction and plunging stock prices¹⁴.

REFERENCES

- [1] Y. Boshmaf, I. Muslukhov, K. Beznosov, and M. Ripeanu, "The socialbot network: When bots socialize for fame and money," in *Proceedings of the 27th Annual Computer Security Applications Conference*, ser. ACSAC '11. New York, NY, USA: ACM, 2011, pp. 93–102. [Online]. Available: <http://doi.acm.org/10.1145/2076732.2076746>
- [2] S. T. Christian Müller-Schloer, *Organic Computing - Technical Systems for Survival in the Real World*. Springer, Berlin, 2018. [Online]. Available: https://www.ebook.de/de/product/29942793/christian_mueller_schloer_sven_tomforde_organic_computing_technical_systems_for_survival_in_the_real_world.html
- [3] S. Tomforde, B. Sick, and C. Müller-Schloer, "Organic computing in the spotlight," *arXiv preprint arXiv:1701.08125*, 2017.
- [4] S. Rudolph and S. Tomforde, "A taxonomy for organic computing systems regarding mutual influences," Institute of Computer Science, University of Augsburg, Tech. Rep., April 2016, last visited 2018-11-28. [Online]. Available: <https://opus.bibliothek.uni-augsburg.de/opus4/frontdoor/index/index/docId/3717>

- [5] M. Tsvetkova, R. Garca-Gavilanes, L. Floridi, and T. Yasserli, "Even good bots fight: The case of wikipedia," *PLOS ONE*, vol. 12, no. 2, pp. 1–13, 02 2017.
- [6] M. Bene, "Go viral on the facebook! interactions between candidates and followers on facebook during the hungarian general election campaign of 2014," *Information, Communication & Society*, vol. 20, no. 4, pp. 513–529, jun 2016.
- [7] Z. Gilani, R. Farahbakhsh, and J. Crowcroft, "Do bots impact twitter activity?" in *Proceedings of the 26th International Conference on World Wide Web Companion*, ser. WWW '17 Companion. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2017, pp. 781–782. [Online]. Available: <https://doi.org/10.1145/3041021.3054255>
- [8] C. Müller-Schloer, C. von der Malsburg, and R. P. Würt, "Organic computing," *Informatik-Spektrum*, vol. 27, no. 4, pp. 332–336, Aug 2004. [Online]. Available: <https://doi.org/10.1007/s00287-004-0409-6>
- [9] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Measurement and analysis of online social networks," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement - IMC '07*, ser. IMC '07. New York, NY, USA: ACM, 2007, pp. 29–42.
- [10] S. Franklin and A. Graesser, "Is it an agent, or just a program?: A taxonomy for autonomous agents," in *Intelligent Agents III Agent Theories, Architectures, and Languages*. Springer Berlin Heidelberg, 1997, pp. 21–35.
- [11] Y. Boshmaf, I. Muslukhov, K. Beznosov, and M. Ripeanu, "Design and analysis of a social botnet," *Computer Networks*, vol. 57, no. 2, pp. 556 – 578, 2013, botnet Activity: Analysis, Detection and Shutdown.
- [12] S. Rudolph, S. Tomforde, B. Sick, and J. Hahner, "A mutual influence detection algorithm for systems with local performance measurement," in *2015 IEEE 9th International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*. IEEE, 2015, pp. 144–149.
- [13] S. C. Woolley, "Automating power: Social bot interference in global politics," *First Monday*, vol. 21, no. 4, mar 2016.
- [14] S. Tomforde, H. Prothmann, J. Branke, J. Hähner, M. Mnif, C. Müller-Schloer, U. Richter, and H. Schmeck, "Observation and control of organic systems," in *Organic Computing A Paradigm Shift for Complex Systems*. Springer, 2011, pp. 325–338.
- [15] M. Forelle, P. N. Howard, A. Monroy-Hernández, and S. Savage, "Political bots and the manipulation of public opinion in venezuela," *CoRR*, vol. abs/1507.07109, 2015. [Online]. Available: <http://arxiv.org/abs/1507.07109>
- [16] V. Subrahmanian, A. Azaria, S. Durst, V. Kagan, A. Galstyan, K. Lerman, L. Zhu, E. Ferrara, A. Flammini, and F. Menczer, "The DARPA twitter bot challenge," *Computer*, vol. 49, no. 6, pp. 38–46, jun 2016.
- [17] C. Freitas, F. Benevenuto, S. Ghosh, and A. Veloso, "Reverse engineering socialbot infiltration strategies in twitter," in *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*, ser. ASONAM '15. New York, NY, USA: ACM, 2015, pp. 25–32. [Online]. Available: <http://doi.acm.org/10.1145/2808797.2809292>

¹⁴<https://www.theguardian.com/business/2015/apr/22/2010-flash-crash-new-york-stock-exchange-unfolded>, accessed on 28.11.18