

Steps toward Understanding the Design and Evaluation Spaces of Bot and Human Knowledge Production Systems

Lei (Nico) Zheng
Stevens Institute of Technology
Hoboken, NJ
lzheng9@stevens.edu

Christopher M. Albano
Stevens Institute of Technology
Hoboken, NJ
calbano@stevens.edu

Jeffrey V. Nickerson
Stevens Institute of Technology
Hoboken, NJ
jnickerson@stevens.edu

ABSTRACT

Bots and humans can combine in multiple ways in the service of knowledge production. Designers make choices about the purpose of the bots, their technical architecture, and their initiative. That is, they decide about functions, mechanisms, and interfaces. Together these dimensions suggest a design space for systems of bots and humans. These systems are evaluated along several criteria. One criterion is productivity. Another is their effects on human editors, especially newcomers. A third is sustainability: how they persist in the face of change. Design and evaluation spaces are described as part of an analysis of Wiki-related bots: two bots and their effects are discussed in detail, and an agenda for further research is suggested.

CCS CONCEPTS

Collaborative and social computing theory, concepts and paradigms • HCI theory, concepts and models • Interactive systems and tools

KEYWORDS

Bots, knowledge production, design space, Wikipedia, Wikidata

1 Introduction

Bots and humans together are responsible for the production of knowledge in Wikipedia. Increasingly, bots are being used to generate data, convert data to articles, and maintain both data and articles. This is part of a more general trend: outside Wikipedia, autonomous tools are being used to generate myriad artifacts, including computer chips and video games [25, 26].

Editors are learning from experience how to use bots and how to design and build bots. Most bots use basic regular expressions or heuristic evaluation, and function as straightforward tools. Yet some bots are, like editors, learning from experience how to clean, classify, and generate data and text. While, in online communities, the bot design process is mainly bottom up, with multiple bots doing similar things, and bots collaborating, there is a governance structure [12]. It would be useful to better understand the design space for these systems. That is, what is possible, now and in the near future? And what is the current coverage of the possible? Such an understanding might be used to discover parts of the design space that might be productively explored by systems designers, bot designers, editors, and bots. It also will be useful to understand the evaluation space, the criteria used to figure out if a particular bot should be deployed. Software designers shuttle

between design space, where they generate variation, and evaluation space, where they figure out which variations deserve further development [21]. Bot designers likely do the same, with the constraint that evaluation involves a governance process.

On Wikipedia, articles are increasingly generated by teams of bots and humans: as early as 2007, most edits were done by bots, although persistent edits were mainly inserted by humans [24]. In 2012, Wikimedia launched Wikidata, which has become a source for bots that create content [30]. Bots continue to evolve, performing more and more complex tasks [5, 23]. In other words, bots are no longer automatic tools that execute precoded functions. Instead, they change the overall ecosystem by their interactions with their operators, administrators, and human Wikipedia editors, as well as other bots. Understanding to what degree and how the bots, the people, and the relationships are changing is important if we want to increase our collective ability to create, integrate, and use knowledge.

The bot-human environment of Wikipedia is an ecosystem with bots and humans playing distinct roles in the creation and maintenance of knowledge. Counter-vandalism is one of the many roles bots play [11, 16, 27]. Recently, researchers have begun to analyze the use of bots in managing Wikidata, a website that mainly maintained by machines rather than human [23]. Humans in some cases seem to bond with the bots as teammates [7–9]. We want to better understand how this affects the overall design of the knowledge production systems and the outcomes of such systems.

2. Design and Evaluation Spaces

2.1 Dimensions of design space

There are many different types of design spaces possible when large numbers of automated agents are interacting. On the one hand are arrangements in which bots are part of strictly delineated hierarchical structures [13]. Such bots are often modular in some way, and share characteristics that make it possible to predict how they will interact. On the other hand are arrangements of bots that are heterogeneous and share little if any common architecture. Wikipedia and Wikidata bots are on this end of the spectrum; they can collaborate, but this collaboration is not forced [12].

On closer look, there is process even in this bottom-up environment. Bots are proposed, their developers questioned, test runs inspected by those outside the team, and discussions entertained before bots are approved for production use.

In attempting to define a design space, we are interested in capturing current variation, in determining factors, so that gaps in a space can be discovered and filled. We are also interested in providing a structure that cuts down the amount of unproductive exploration that might be done. That is, we want to find some dimensions or choice points that are generative of useful solutions.

With respect to bots, we are interested in the design of human bot systems. We have identified three sets of choices that need to be made by designers: this can be considered the start of characterization of the design space.

The first has to do with the function of the bot: what its purpose is, and, in particular, which types of pages or data it operates on. Some bots post into articles, others onto user talk pages, and still others create or modify tuples in the Wikidata knowledge graph. The second has to do with the technical architecture of the bot: how it is coded, and if and how it shares function libraries or other forms of embedded technology.

The third is the extent and nature of the interface with humans and in particular, the initiative that can be taken by both. Humans can invoke bots. Bots can revoke edits. Humans can monitor bot edits, and bots can monitor human edits. Humans can monitor the bots that monitor humans. That is, there is a recursive structure in the relationship between humans and bots, and designers anticipate this structure. Moreover, the relationship can function at different levels. Specifically, Wikipedia involves a kind of implicit coordination at the group level that manifests in Wikiprojects [23]. Bots can service group as well as individual goals

2.2 Evaluating bot and human systems

Bots are created to aid humans in performing certain tasks. They can help in some ways but hinder in others: the evaluation space spells out the criteria to be used in evaluating the design of a bot and human system.

Bots support human work by improving productivity [10, 28, 31]. For any technology, the overall impact on productivity is often a major criterion that drives adoption [3]. But sometimes bots send misleading messages, take disruptive actions and increase the maintenance burden on humans [28, 31]. Moreover, some of the bots' impacts are hard for developers to anticipate [1, 15]. It is clear bot evaluation involves more than measuring productivity.

In Wikipedia, the interactions between editors are mediated through articles and talk pages. These interactions can be measured with respect to frequency. That is, the overall number of interactions between editors, as well as the intervals between edits, may influence further decisions to edit [32]. The interactions may express attitudes that may trigger emotional responses. For example, articles for deletion pages record debates that often turn emotional. Just as product designers pay attention to the emotional reactions of users, those evaluating bots pay attention to the reactions of human editors. In particular, due to the importance of recruiting volunteer labor in the form of editors

in Wikipedia, encouraging rather than discouraging newcomers is important. This can be measured by looking at newcomer survival rates.

2.3 How bots affect newcomers' survival

As an illustrative example, we explore the effect of two distinct Wikipedia bots on newcomers' survival rates: the ClueBot NG and the SuggestBot.

ClueBot NG is the anti-vandal bot that uses machine learning algorithms to detect vandalism on Wikipedia content pages [29]. After removing an instance of suspected vandalism, it will leave a message on the contributor's user talk page to inform the editor about the action taken. The bot was formally launched on the entire English Wikipedia in October 2010 and has reverted over 5 million instances of suspected vandalism. Even though it has successfully defended the encyclopedia from vandalism, Halfaker et al. showed that the rude greeting this bot provided to those practicing their first edits decreased retention of newcomers [15].

SuggestBot is a recommender bot that uses collaborative filtering to help users find articles to contribute [6]. It delivers a list of suggested articles to its subscribers based on their recent edit history. The bot was launched in December 2005 and has provided over 64,000 suggestions to 7,843 editors.

We are interested in how communication with bots affects newcomers' retention. It is reasonable to expect that the newcomers who received a friendly suggestion will have a higher engagement level and thus participate longer on the platform. By contrast, newcomers who receive a warning message about a recent edit that was reverted are more likely to feel offended and leave the platform.

To better understand this, we analyzed all the newcomers that registered in 2011 and made at least one edit in that year. From the newcomers we selected all the users who received a message from either of the two bots discussed above within 60 days of their first edit and separate them into the two groups. We also randomly selected three thousand newcomers who did not receive any messages as the comparison group. We define a newcomer "active" in a certain month if the newcomer made at least one edit in that month. We tracked newcomer activity in the following two years and calculated monthly group survival rates based on how many editors were active in that month. Table 1 shows the summary statistics of newcomers in each user group.

User Group	Number of Editors	Retention rate
Suggested Newcomers	92	23.73%
Reverted Newcomers	34,144	0.60%
Random Newcomers	3,000	1.75%

Table 1: Summary statistics of three user groups.

Similar to Halfaker’s finding, newcomers who were reverted by the ClueBot NG have a survival rate lower than the random newcomers. However, newcomers that were greeted by the SuggestBot have a significantly higher survival rate, 13 times larger than the random newcomers. Because newcomers have to actively register for SuggestBot, one possible explanation is that only the most motivated editors register in the first place. The degree to which retention is due to self-selection or to the effects of positive interaction can’t be shown from this study. Most likely both factors are at play; a random assignment experiment would be one way to further explore the issue. Figure 1 shows the monthly survival rate of the three groups.

We also find that another potential evaluation criterion, the frequency of bot-human interactions, has an effect. Figure 2 shows the monthly survival rate of newcomers who received suggestions once and who received suggestions more than once. As we can see from the figure, newcomers who received more suggestions generally will have a greater survival rate than newcomers who only received one suggestion, although the advantage diminishes over time. This may be because of natural attrition in editors, or it may be the result of habituation. Either way, it suggests that evaluation criteria be measured over sufficiently long periods of time to capture such trends.

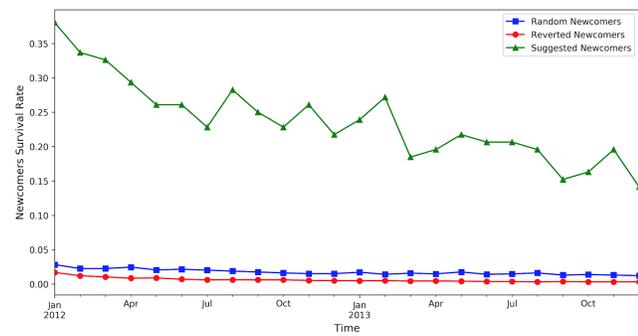


Figure 1: The monthly survival rate of newcomers.

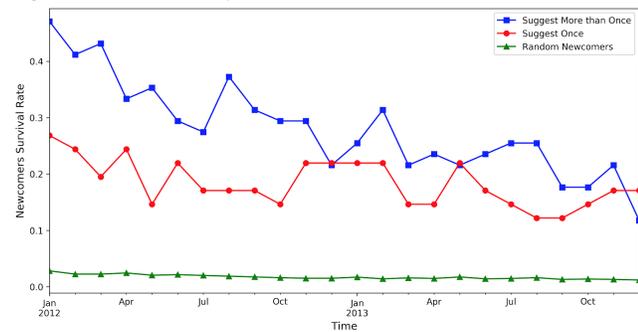


Figure 2: The monthly survival rate of suggested newcomers.

2.4 Relationships between bots and developers

Software is not only difficult to build, it is also difficult to maintain, and, in the software engineering world, maintainability is an important evaluation criterion. To better understand this, we looked more deeply into the construction and management of the two bots. The users that created SuggestBot are ForteTuba and Nettrom. Both of them are academics, and declare their public

identities on their user pages. User ForteTuba began editing on Wikipedia on October 24, 2005. An archived discussion thread comment by ForteTuba reads “Just an update, SB has made suggestions to a couple hundred people unsolicited without complaints, and to a few dozen who asked for recommendations. Overall feedback has been positive on its talk page.” Implicit in this comment is a process of allowing the bot to be evaluated in its early stages of development. Bots need to gain approval from the overall Wiki community before being deployed in a process mediated through discussions on a page devoted to Bot requests for approval (BRFA). For BRFA#2, ForteTuba took the lead and was listed as the bot operator on the request for approval, saying he would make a “substantial rewrite to fix known flaws and generate better recommendations.” He also mentions the purpose of the bot, saying “I think it can be used to improve the feelings of community in Wikipedia by pointing people to other people and groups, so I intend to do that.”

However, when BRFA#3 was proposed for SuggestBot, a new operator was listed: username Nettrom, showing a transition in ownership of the bot. The purpose of this new request for approval was “As part of a research study, we want to post suggestions to a limited number of new users in order to see if we can make them contribute more than they otherwise would.” BRFA #3 took place in 2010 as part of an experiment. BRFA #5 was also related to research. However, Nettrom also made BRFA #4, in which “we now want to automate the process of handling our users who have signed up to receive suggestions regularly to ease the burden of having to do this manually.” This task makes sense to ease the burden of manual editing, but Nettrom is demonstrated as the clear owner of the bot after BRFA#3. To summarize the other BRFAs, BRFA#6 is used for one time requests, BRFA#7 was made to update the community portal’s list of open tasks. BRFA#8 was made to update various task templates periodically.

The interesting phenomenon is the clear transition in ownership between ForteTuba to Nettrom. This is an example of sustainability: while bots can extend the agencies of individual editors, they also can, like many technologies, be transferred to others. In other words, bots can be general purpose tools that can be used by others. Ownership can be shared, and responsibility can transfer. This increases the potential longevity of the bot, and may also increase its maintainability: designers may anticipate handing over control to make such handovers simpler.

During this transfer of responsibility, SuggestBot did not lose any productivity in terms of edit count. Both operators remain constantly engaged during their ownership (see Figure 3). It still remains impactful on the users of Wikipedia that choose to use it. While perhaps this seamless transfer is not surprising because these two operators were colleagues, it points to the general issue of the provenance, ownership, and responsibilities related to bots. Bots may outlive their designers, or bots may be reinvented every few years. Bots might eventually design other bots.

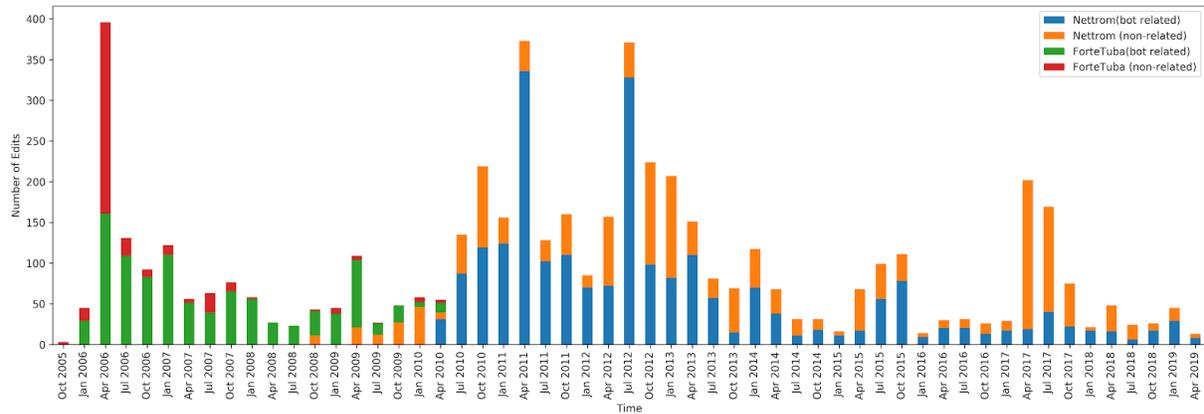


Figure 3: Bot transition of ownership

2.5 Bots and their governors

While editors commend bots for streamlining the knowledge production process, they also complain that bots not only solve problems, but also create them [31]. On the path of making bots more effective, many complex machine learning problems need to be solved. For example, to build a bot that protects pages, we may want to use a machine learning classifier to identify whether there is a possible violation. Then we will want to add another classifier to indicate which policy was likely to be violated; then in order to let the bot speak properly, we may want to add a Natural Language Generation module; if we also want the bot to learn from past experience, we should apply online learning models that update the algorithm itself. Finally, we should train another classifier to examine the bot and the training data so we know that the bot is not exhibiting some form of bias. In other words, the design space of Wikipedia bots may include the currently understood variants of applied machine learning systems.

In a variety of online communities, bots are created in a bottom-up approach by enthusiastic volunteers. In such environments, a developer may create a bot without knowing another bot already handles a similar job. A bot may not be able to share data and resources with other bots. They cannot learn from others' experience.

A bot may easily fall into a dormant state because its developer is no longer active. In this case, another developer may notice and write another similar bot. The original bot may be abandoned. From a higher level, the bots are evolving the way software evolves, with some bots modified and improved while others are eventually culled. These issues put the governance of bots in the spotlight. Currently the Wikipedia community has a Bot Approvals Group (BAG) to govern bot approvals and solve related disputes.

3. What ought we to do

3.1 Steps for understanding the design space

Bots in Wikipedia have a variety of functions, including collecting information, executing actions and generating content [30]. As of February 2019, there were 1,939 registered bot accounts and many of them interacted with the human editors.

With respect to function, bot designers declare what they do when they apply for approval and again on the bot user pages when approved. These declarations might be enumerated to document the range of current functionality. This can be done manually, although extracting functional descriptions automatically is a current area of research [18]. These functions might also be analyzed from a time series perspective to better understand what functions are likely to be developed in the future. For example, relatively few bots currently focus on Wikidata, but a slate of recent developments suggest that more such bots can be expected. The challenge is to anticipate, to sketch out what can conceivably be built.

Bot architectures can be examined in cases where bot owners make source code available. Many different languages and frameworks are in use, and each developer makes a design choice to use a particular toolkit. We can imagine the design of tighter specifications for bots that might encourage more uniformity, which might aid recombinations in the design phase and systematic monitoring in the management phase. Given the nature of the Wikimedia community, it is unclear if or how this will happen.

With respect to interfaces, advances in machine learning are leading to machines that can operate more autonomously, and hence can take initiative [14, 17, 19]. The design space issue is the following. Most bots can take initiative, and can likewise allow humans to initiate. The challenges lie in the interactions over time. A dispute between humans and bots over reversion of content can easily escalate, just as disputes between human editors can. Creating productive interactions between bots and humans involves pulling from not only the design space of machine learning, but also psychology. Understanding the valence and arousal of human interactions is a challenge for bots, and the next step is even more challenging: figuring out when to defuse or escalate. Bot designers may want to create

experimental bots that test the extremes of design space, but other community members may be resistant to training errant bots at the expense of producing more knowledge.

3.2 Steps for understanding the evaluation space

Current evaluation practices can be understood by looking at the bot requests for approvals and comments in response. This might allow an enumeration of different criteria.

Such criteria will involve tradeoffs. It may be desirable to get a bot deployed quickly, but that bot may incur load on human editors who will end up fixing the bot's mistakes. The bot approval process sometimes discusses these tradeoffs.

Bots certainly can affect editing patterns in a negative way: reverting newcomer's edits can decrease retention. Bots may affect editing patterns in a positive way if the tone of interaction is positive, and the speech act is generative: suggestions seem to be well received.

Experimental rather than observational studies may be necessary to understand this better. To avoid self-selection biases (for example, perhaps only highly motivated editors sign up for suggestbot), random assignment to experimental conditions could be used. On the other hand, environments such as Wikipedia are production environments with social goals that may not allow long running blinded experiments.

One alternative is to perform experiments in specialized Wikis that are not in production. In addition, observational data can continue to be collected to better understand what impacts the interactions between bots and humans have on editing behaviors and on the quality of knowledge produced.

Several issues may admit to such exploration. Frequency of interaction may itself be influencing retention [33], perhaps interacting with other factors. One potential factor is the sentiment of the message. A second is the speech act of the message: for example, whether the bot is making suggestions, asking questions, or issuing direct imperatives [32].

With respect to evaluating bots, there are two perspectives to consider. We discussed the effects a bot has on editors. But there is another side; editors design bots and then work with these bots to develop content. This kind of collaboration is often increasing editors' productivity by large factors. Moreover, editors may be building bots that support a kind of supply chain of knowledge, with facts entered into Wikidata, cleaned by bots, then used by human bot teams to generate infoboxes and portions of articles. This hybrid collaboration may create a sense of entitativity; a feeling of belonging to a team of humans and bots [2, 20]. The extent to which this leads to a communal sense of knowledge generation between larger sets of humans and bots is worth considering. For some editors, bots may provide a kind of extended agency; for others, bots may be perceived as team members. Measuring these differences may be important for understanding what drives the further development of bots.

3.3. Concluding thoughts

It will be useful to understand how the process of designing, deploying and running bots changes over time, and how the activities and skills of the humans change as the bots change. Wikipedia provides traces of machine-human interaction that might provide clues as to how increasingly autonomous machine agents change the nature of human work and the nature of knowledge artifacts.

ACKNOWLEDGMENTS

This research is based on work supported by the National Science Foundation under grants IIS-1422066, CCF-1442840, IIS-1717473, and IIS-1745463.

REFERENCES

- [1] Angwin, J., Larson, J., Mattu, S. and Kirchner, L. 2016. Machine bias: There's software used across the country to predict future criminals. *And it's biased against blacks*. *ProPublica*.
- [2] Blanchard, A.L., Caudill, L.E. and Walker, L.S. 2018. Developing an entitativity measure and distinguishing it from antecedents and outcomes within online and face-to-face groups. *Group processes & intergroup relations: GPIR*.
- [3] Brynjolfsson, E. 1993. The Productivity Paradox of Information Technology. *Communications of the ACM*. 36, 12 (Dec. 1993), 66–77.
- [4] Chan, J., Fu, K., Schunn, C., Cagan, J., Wood, K. and Kotovsky, K. 2011. On the Benefits and Pitfalls of Analogies for Innovative Design: Ideation Performance Based on Analogical Distance, Commonness, and Modality of Examples. *Journal of Mechanical Design*. 133, 8 (Aug. 2011), 081004.
- [5] Chisholm, A., Radford, W. and Hachey, B. 2017. Learning to generate one-sentence biographies from Wikidata. *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*
- [6] Cosley, D., Frankowski, D., Terveen, L. and Riedl, J. 2007. SuggestBot: Using Intelligent Task Routing to Help People Find Work in Wikipedia. *Proceedings of the 12th International Conference on Intelligent User Interfaces (New York, NY, USA, 2007)*, 32–41.
- [7] Geiger, R.S. 2013. Are computers merely supporting cooperative work: towards an ethnography of bot development. *Proceedings of the 2013 conference on Computer supported cooperative work companion (2013)*, 51–56.
- [8] Geiger, R.S. 2014. Bots, bespoke, code and the materiality of software platforms. *Information, Communication and Society*. 17, 3 (Mar. 2014), 342–356.
- [9] Geiger, R.S. 2011. *The Lives of Bots*. (2011).
- [10] Geiger, R.S. 2009. The Social Roles of Bots and Assisted Editing Programs. *Proceedings of the 5th International Symposium on Wikis and Open Collaboration (New York, NY, USA, 2009)*, 30:1–30:2.
- [11] Geiger, R.S. and Halfaker, A. 2013. When the Levee Breaks: Without Bots, What Happens to Wikipedia's Quality Control Processes? *Proceedings of the 9th International Symposium on Open Collaboration (New York, NY, USA, 2013)*, 6:1–6:6.
- [12] Geiger, R. S. and Halfaker, A. 2017. Operationalizing Conflict and Cooperation between Automated Software Agents in Wikipedia: A Replication and Expansion of 'Even Good Bots Fight'. *Proc. ACM Human Computer Interaction* 1, CSCW, Article 49 DOI: <https://doi.org/10.1145/3134684>
- [13] Gerkey, B.P. and Mataric, M.J. 2004. A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems. *The International journal of robotics research*. 23, 9 (Sep. 2004), 939–954.
- [14] Gombolay, M., Bair, A., Huang, C. and Shah, J. 2017. Computational design of mixed-initiative human-robot teaming that considers human factors: situational awareness, workload, and workflow preferences. *The International journal of robotics research*.
- [15] Halfaker, A., Geiger, R.S., Morgan, J.T. and Riedl, J. 2013. The rise and decline of an open collaboration system: How Wikipedia's reaction to popularity is causing its decline. *The American behavioral scientist*. 57, 5 (2013), 664–688.
- [16] Halfaker, A. and Riedl, J. 2012. Bots and cyborgs: Wikipedia's immune system. *Computer*. 45, 3 (2012), 79–82.

- [17] Hoffmann, R., Amershi, S., Patel, K., Wu, F., Fogarty, J. and Weld, D.S. 2009. Amplifying community content creation with mixed initiative information extraction. *Proc. CHI '09* (2009).
- [18] Hope, T., Chan, J., Kittur, A., and Shahaf, D. 2017. Accelerating Innovation Through Analogy Mining. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*.
- [19] Liapis, A. and Yannakakis, G.N. 2016. Boosting computational creativity with human interaction in mixed-initiative co-creation tasks. *Proceedings of the ICCG workshop on Computational Creativity and Games* (2016).
- [20] Nash, K., Lea, J.M., Davies, T. and Yogeewaran, K. 2018. The bionic blues: Robot rejection lowers self-esteem. *Computers in human behavior*. 78, (Jan. 2018), 59–63.
- [21] Nickerson, J.V. and Yu, L. 2013. *Going Meta. Software Designers in Action: A Human-Centric Look at Design Work*. Chapman and Hall/CRC. 321–342.
- [22] Ozturk, P. and Nickerson, J. 2015. Paths from Talk to Action. *Proceedings of the International Conference on Information Systems*.
- [23] Piscopo, A. and Simperl, E. 2018. Who Models the World?: Collaborative Ontology Creation and User Roles in Wikidata. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 141
- [24] Priedhorsky, R., Chen, J., Lam, S. (tony) K., Panciera, K., Terveen, L. and Riedl, J. 2007. Creating, Destroying, and Restoring Value in Wikipedia. *Proceedings of the 2007 International ACM Conference on Supporting Group Work* (New York, NY, USA, 2007), 259–268.
- [25] Seidel, S., Berente, N., Lindberg, A., Lyytinen, K. and Nickerson, J.V. 2019. Autonomous Tools and Design: A Triple-loop Approach to Human-machine Learning. *Communications of the ACM*. 62, 1 (2019), 50–57.
- [26] Seidel, S., Berente, N., Martinez, B., Lindberg, A., Lyytinen, K. and Nickerson, J.V. 2018. Autonomous Tools in System Design: Reflective Practice in UbiSofts Ghost Recon Wildlands Project. *Computer*. 51, 10 (Oct. 2018), 16–23.
- [27] Steiner, T. 2014. Bots vs. Wikipedians, Anons vs. Logged-Ins (Redux): A Global Study of Edit Activity on Wikipedia and Wikidata. *Proceedings of The International Symposium on Open Collaboration* (New York, NY, USA, 2014), 25:1–25:7.
- [28] Storey, M.-A. and Zagalsky, A. 2016. Disrupting Developer Productivity One Bot at a Time. *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering* (New York, NY, USA, 2016), 928–931.
- [29] Stvilia, B., Twidale, M.B., Smith, L.C. and Gasser, L. 2005. Assessing Information Quality of a Community-Based Encyclopedia. *ICIQ*. 5, (2005), 442–454.
- [30] Vrandečić, D. and Krötzsch, M. 2014. Wikidata: A Free Collaborative Knowledgebase. *Communications of the ACM*. 57, 10 (Sep. 2014), 78–85.
- [31] Wessel, M., de Souza, B.M., Steinmacher, I., Wiese, I.S., Polato, I., Chaves, A.P. and Gerosa, M.A. 2018. The Power of Bots: Understanding Bots in OSS Projects. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 182
- [32] Xiao, L. and Nickerson, J. 2019. Imperatives in Past Online Discussions: Another Helpful Source for Community Newcomers? *Proceedings of the 52nd Hawaii International Conference on System Sciences*.
- [33] Zheng, L., Mai, F., Gordon, D. M., and Nickerson, J. V. 2018. Bursty Coordination in Online Communities. *Workshop on Information in Technology*.