

Evaluating the impact of the Wikipedia Teahouse on newcomer socialization and retention

Jonathan T. Morgan
Wikimedia Foundation
jmorgan@wikimedia.org

Aaron Halfaker
Wikimedia Foundation
ahalfaker@wikimedia.org

ABSTRACT

Effective socialization of new contributors is vital for the long-term sustainability of open collaboration projects. Previous research has identified many common barriers to participation. However, few interventions employed to increase newcomer retention over the long term by improving aspects of the onboarding experience have demonstrated success. This study presents an evaluation of the impact of one such intervention, the Wikipedia Teahouse, on new editor survival. In a controlled experiment, we find that new editors invited to the Teahouse are retained at a higher rate than editors who do not receive an invite. The effect is observed for both low- and high-activity newcomers, and for both short- and long-term survival.

CCS CONCEPTS

•Human-centered computing → Empirical studies in HCI; Empirical studies in collaborative and social computing;

KEYWORDS

open collaboration; online communities; Wikipedia; virtual work; social norms; socialization

ACM Reference format:

Jonathan T. Morgan and Aaron Halfaker. 2018. Evaluating the impact of the Wikipedia Teahouse on newcomer socialization and retention. In *Proceedings of The 14th International Symposium on Open Collaboration, Paris, France, August 22–24, 2018 (OpenSym '18)*, 7 pages. DOI: 10.1145/3233391.3233544

1 INTRODUCTION

Open online collaborations can only thrive if they are able to recruit a steady stream of newcomers to replace the people who stop contributing. New contributors to open collaboration projects like Wikipedia and open source software face both social and technical barriers that make it difficult for them to learn how to participate productively, and discourage them from doing so [11, 28, 31].

The sociotechnical nature of these barriers make it difficult to design interventions to fully and effectively socialize newcomers. Mechanisms for new contributor onboarding on a large, mature open collaboration like English Wikipedia¹ must take into account

¹All references to Wikipedia refer to the English language edition, unless otherwise stated.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

OpenSym '18, Paris, France

© 2018 Copyright held by the owner/author(s). 978-1-4503-5936-8/18/08...\$15.00
DOI: 10.1145/3233391.3233544

both the technical challenges that newcomers face (e.g., learning the markup syntax used to format articles), and the social challenges (e.g., navigating a complex bureaucracy of written and unwritten rules [2, 18]). In addition, new Wikipedia editors often face mistrust and even outright hostility from more experienced editors who view their inexperience and unknown motivations as a threat to the quality of Wikipedia, and who have embedded these assumptions into the tools they use to interact with newcomers[12]. Finally, the voluntary and transient nature of contribution to open collaborations requires that mechanisms for new editor socialization scale well to accommodate a high volume of potential new long-term contributors, even if only a relatively small number of those new arrivals are likely to stick around [7].

Wikipedia editors and the Wikimedia Foundation have tried a variety of approaches to address the problem of new editor socialization and retention over the past 15 years. The pace of experimentation around new editor onboarding has accelerated since 2009, when researchers first identified a decline in the population of active contributors, which had grown exponentially until 2007 [30]. This editor decline has been shown to be due in part to a steady increase in hostility towards good faith newcomers—those who join Wikipedia to contribute productively, not vandalize—over time, which led to a concomitant decrease in the rate at which good faith new editors were retained as long-term contributors [11].

Approaches to onboarding have ranged from the purely social (e.g. individual mentorship of newcomers by patient Wikipedia veterans) to purely technical (e.g. software features intended to automate aspects of the socialization process without direct human intervention). So far, neither purely social or purely technical efforts have been shown to be effective at providing effective socialization at a scale that leads to a substantial increase in the number of new editors who go on to become Wikipedians.

In this paper, we summarize the state of research and design around new editor socialization on Wikipedia, and analyze one of the most potentially impactful retention mechanisms that has been attempted, a new editor support forum called the Teahouse. We perform the first controlled study to determine the impact of being invited to the Teahouse on a new editor's likelihood of continued contribution over the short and long term. We show that new editors who are invited to the Teahouse are significantly more likely to continue contributing after three weeks, two months, and six months than a similar cohort who were not invited.

To our knowledge, this is the first time any intervention has shown positive impact on Wikipedia editor retention in a controlled study. We discuss the implications of our findings for both the design and evaluation of sociotechnical onboarding interventions, and identify specific areas for further research and development.

2 BACKGROUND

Many researchers have examined the challenges associated with newcomer socialization across a range of online communities, include Usenet newsgroups [3], social news websites [14], online forums [26], scientific collaborations [1], and open source software projects [6].

While gaining new contributors is vital to long-term community health, it can be a source of friction in the short term. Newcomers frequently and unintentionally violate community norms, disrupt community activities, all of which can reduce existing community members' motivation to participate [14]. A large influx of newcomers into Usenet newsgroups in the early 1990s prompted existing community members to complain that their community was now experiencing an "Eternal September"², a reference to the annual increase in new registrations that usually came at the beginning of the North American academic year. Wikipedia's own "Eternal September" began around 2006, as increasing visibility of Wikipedia led to an exponential increase in new editors who were eager to contribute to the still relatively young encyclopedia, as well as an increasing number of 'vandals' who wanted to disrupt or co-opt Wikipedia.

2.1 Community socialization strategies

In response to this influx, the Wikipedia community overall became less willing to discriminate between well-intentioned newbie mistakes and authentic vandalism. Editors who monitored new contributors and contributions chose to focus their efforts on maintaining the quality of the encyclopedia at the cost of long term health and sustainability. As a result, Wikipedia began to experience a decline in new editor retention that shrunk the community even as the number of articles that required monitoring and maintenance continued to grow[11].

New quality control technologies developed by Wikipedia community members, such as autonomous anti-vandalism bots and specialized user interfaces for efficiently reviewing and reverting incoming edits and new articles, were highly effective at minimizing potential damage by prioritizing recall and speed over precision and care [12]. However, community-driven approaches to socializing this influx of newcomers were less effective. Generic warning messages delivered—often by bots—to new editors whose contributions were rejected did not dissuade vandals, but did dissuade good faith newcomers [9]. Generic welcome messages—long lists of links to policies, how-tos, and help documentation—delivered the same way were equally ineffective [34].

Social programs such as Adopt-a-User [19] and invitations to join wiki-projects [4] provided more personalized socialization opportunities. Although these purely social initiatives showed some promise, they required a substantial level of individual effort to maintain, and could therefore only be made available to a relatively small number of relative newcomers—generally, to those who were already highly likely to be retained because they had overcome many of the initial social and technical barriers and demonstrated a willingness to continue editing for several weeks or months [21].

2.2 Wikimedia Foundation interventions

Starting in 2011, the Wikimedia Foundation began to perform its own research and development activities focused on understanding and addressing the causes of new editor attrition³.

In April of 2012, the Wikimedia Foundation formed the Growth Team⁴, a product development team that was specifically tasked with addressing the decline in newcomer retention through intervention and experimentation on top of the MediaWiki platform. They documented their intentions, methods, and results in a series of white papers published on the Wikimedia Meta-wiki. In their experiments, they focused on minor changes to the software that might support motivation or task routing for new editors. For example, they experimented with reporting milestones to new editors (e.g. "You've saved your 10th edit!") [22], they experimented with making the account creation process easier [32], and they experimented with providing new editors with in-context editing instructions and task suggestions [10]. While the team experimented with a variety of strategies (task routing, in-context help, motivational, etc.) and points of intervention (editing anonymously, registration process, immediately post-registration, etc.), no study showed a significant effect new editor retention rates.

By October of 2014, the newcomer intervention experimentation efforts were abandoned and the members of the Growth Team were reassigned to other work in the organization. While the Growth team's exploration was not exhaustive, it shows a clear indication that a large set of obvious, technical interventions are ineffective in addressing Wikipedia's newcomer retention issues. This suggests that the Wikipedia's newcomer retention is difficult to solve and that the solution might not lay in the space of naive user-interface interventions.

Other Wikimedia Foundation experiments in socialization combined both technological and social components to facilitate newcomer socialization on a large scale while preserving a degree of personalization and addressing a broader set of socialization challenges.

The MoodBar tool combined a UI widget through which new editors could request in-context help and a feed interface to help experienced editors triage and respond to a large number of incoming help requests. An evaluation of MoodBar suggested a small but consistent increase in retention between editors who were provided with the intervention and users who were not [5]. However, because the treatment and control groups were collected at two different time periods, it was not possible to determine whether the effect shown was due to the intervention, or to exogenous factors (e.g. seasonal variations). Experiments with more personalized welcome and warning messages to new editors demonstrated that changing the style of a message could have short-term impacts on subsequent activities, but similarly failed to show any reliable, positive effects on retention [9].

In 2013, Wikimedia developed and piloted The Wikipedia Adventure a gamified new editor training system. Although fully automated, the Wikipedia Adventure was designed to provide both comprehensive socialization and an impression of personalization through the use of self-guided interactive tutorials, realistic training

²https://en.wikipedia.org/wiki/Eternal_September

³https://meta.wikimedia.org/wiki/Research:Wikimedia_Summer_of_Research_2011

⁴<https://www.mediawiki.org/wiki/Growth>

scenarios, levels and badges, and a light, playful tone. Although the Wikipedia Adventure was widely used and rated favorably by the editors who played along, a controlled study of the system did not show any long-term impact [20]. The evaluation of the Wikipedia Adventure is notable because it was the first study of new editor onboarding on Wikipedia to combine an analysis of long-term retention impacts with a randomized controlled study design—a model we follow in our current study.

The only other major newcomer socialization intervention developed by the Wikimedia Foundation, the Wikipedia Teahouse, was created in 2012 and is the focus on the current study. The Teahouse combined social and technical components to provide comprehensive socialization on a large scale, and was designed to promote long-term retention.

2.3 The Wikipedia Teahouse

The Teahouse is a space on Wikipedia designed for new editors, called guests, and staffed by experienced editors, called hosts. The Teahouse contains a Q&A forum where guests can ask questions, and display board where guests can post personal profiles to introduce themselves to the Wikipedia community and view the profiles of other guests and hosts. See Morgan et al. [17] for a more complete description of the Teahouse.

Like MoodBar and The Wikipedia Adventure, the Teahouse was widely used and received positive reviews [16]. Within the first six months, over 1000 new editors had asked questions or created profiles. A preliminary investigation of the impact of Teahouse participation on subsequent editing activity during the first six months performed by Morgan et al [17] showed that, on average, Teahouse guests made more edits to more articles and edited two weeks longer than invitees who chose not to participate.

Although these findings on participant engagement, satisfaction, and subsequent editing activity provide compelling evidence that the Teahouse has the potential to be an effective mechanism for new editor socialization, the design of the retention analysis performed by Morgan et al. was not adequate to determine a causal effect of the Teahouse on retention. The study compared the subsequent editing behavior of invitees who participated in the Teahouse with invitees who did not. It is likely that some invitees who did not visit the Teahouse had already left Wikipedia due to boredom or frustration before they received an invite, and therefore did not have the opportunity to avail themselves of the Teahouse. Furthermore, the study design contains a self-selection bias: invited editors made a conscious decision to undergo the ‘treatment’, and were therefore likely more intrinsically motivated to continue editing than those who declined the invitation.

A different approach is necessary in order to evaluate whether the Teahouse has a positive impact on long-term new editor retention. Fortunately, such an experiment is possible, even many years later, because unlike the Wikimedia Foundation’s other onboarding interventions, the Teahouse was never disbanded or discontinued. Although Wikimedia Foundation staff and contractors led the design and launch of the Teahouse, the forum itself is run by volunteers. The relatively minimal technical requirements of the Teahouse—a bot to send invites and update pages and a Javascript UI

widget to make it easier for newcomers to ask and answer questions—were designed to be easily maintained by volunteers without Wikimedia Foundation support or oversight. As a result, the Teahouse has been active for more than 6 years. As of December 2015, over 8,000 new editors had participated.

This research paper reports on findings from a controlled study in which the experimental manipulation was whether a new editor received a message inviting them to the Teahouse. Previous interview and survey research demonstrated that the Teahouse provides new editors with a welcoming environment, valuable opportunities for socialization, and a sense of belonging that motivates them to continue editing.

Therefore, we hypothesize that new editors who receive an invitation to participate in the Teahouse will be more likely to continue editing Wikipedia for a longer period of time than those who are not invited to participate—even if they do not choose to ask a question or create a profile.

H1: *Compared to a control group, new editors who receive a Teahouse invite are more likely to be retained than those who are not.*

Onboarding interventions may have differential impacts on short- and long-term retention; a system is focused on helping newcomers overcome initial technical hurdles may help these new editors achieve their immediate goals, but not provide the breadth or depth of support necessary for promoting long-term retention. Previous research suggests that individuals who are still contributing after two months are very likely to continue participating for many more months or even years [21]. Because the Teahouse is intended to provide new editors with socialization and support on an ongoing, as-needed basis across a wide variety of topics, we measure retention at several intervals: after 3-4 weeks, 1-2 month, and 2-6 months.

As is the case with other open collaboration systems, the participation rate of both new and established Wikipedia editors is highly variable [33]. Some regular contributors average hundreds, or even thousands of edits per month, while others contribute at a consistent, but much lower, rate. Furthermore, an individual editor’s participation rate may vary month-to-month. In order to measure retention of both lower and higher-volume contributors, we measure survival with both 1 and 5 edits per interval.

3 METHODS

While we wanted to understand the effect that the Teahouse has on new editor retention in Wikipedia, we couldn’t fairly compare editors who took advantage of the Teahouse with those who would have had they been invited. We realized that the only point of intervention that we had control over was point at which editors were invited to the Teahouse. Given that there are no direct links to the Teahouse from articles and few direct links from article discussion pages (the pages newcomers are most likely to visit), it is unlikely that a new editor would find the Teahouse on their own through browsing. Thus an invitation to the Teahouse represents a clear experimental intervention and a non-invitation represents a control. By using these two conditions for comparison, we’re effectively exploring the effect of the invitation with the implication that the greatest value to the newcomer experience happens after a newcomer follows the invitation and learns about the Teahouse.



Hi **Killerbeez3!** Thanks for contributing to Wikipedia. Be our guest at [the Teahouse!](#) The Teahouse is a friendly space where new editors can ask questions about contributing to Wikipedia and get help from peers and experienced editors. I hope to see you there! [I JethroBT \(I'm a Teahouse host\)](#)

[Visit the Teahouse](#)

This message was delivered automatically by your robot friend, [HostBot \(talk\)](#) 17:25, 24 November 2015 (UTC)

Figure 1: An invitation to the Teahouse

3.1 Experimental implementation

Invitations to the Teahouse (Figure 1) are generally delivered by a robot called HostBot, which is maintained by the first author. HostBot periodically queries a database of recently registered new editors to send invitations to. New editors who have made at least 5 edits, have not received serious vandalism warnings, and have not been blocked from editing are sent an invitation. Each invitation is signed by a Teahouse host who had previously agreed to have their name listed on these invitations.

For this study, we split the group of editors who would have normally been sent an invitation into an experimental group (invited) and a control group (noted, but not invited). A small subset of Teahouse invitations are also posted manually by human editors. We excluded newcomers who had received a manual invitation from our analysis.

During the 71 days between October 20, 2014 and January 12, 2015 we sampled users into the control and experimental groups. In each day of the sampling period, 50 editors who met HostBot’s criteria were randomly selected as a control group and invitations were not sent to these users. We kept the control group capped in this way because the findings from [17] suggest that Teahouse had a positive effect for newcomers and we wanted to ensure that we minimized the number of new editors who were disadvantaged by our experiment. The final control group contained a total of 3,092 users. The remaining 11,674 users in the sample received an invitation to the Teahouse from HostBot on their Wikipedia user talkpage.

3.2 Measuring retention

In order to measure the effects that the Teahouse invitation had on retention, we borrowed a retention measurement developed by the Research team at the Wikimedia Foundation called “surviving new editors” [25]. This metric has two terms that can be used to adjust the duration of retention that is being measured: t_2 or the “trial period” and t_3 or the “survival period”. There is also a term m for the number of edits necessary to signal “survival”. An editor is considered surviving when there is evidence that the editor was

retained during the survival period. By adjusting the trial period, the term of retention can be adjusted. For example, our short-term metric used $t_2 = 3$ weeks and $t_3 = 1$ week which means that if an editor saved an edit between 3 and 4 weeks after they register, they’d be considered “surviving”. To explore mid- and long-term retention effects, we also took measurements at $t_2 = 1$ month and $t_3 = 1$ month as well as $t_2 = 2$ months and $t_3 = 4$ months.

Gathering this type of long-term data is difficult because it required us to wait over 6 months to complete our analysis, but it is essential for understanding whether or not the Teahouse invitations had the intended effect on long-term retention. We also compare values of m for 1 edit and 5 edits which represent whether an editor was nominally or substantially active.

To look for differences between the experimental and control groups, we used a χ^2 test to compare proportions of editors deemed “surviving” at each time scale. Our specific hypotheses about retention, grouped by minimum activity threshold, are:

H1a: *Invitees are more likely to have made 1 or more edits 3-4 weeks, 1-2 months, and 2-6 months after the invite date than control.*

H1b: *Invitees are more likely to have made 5 or more edits 3-4 weeks, 1-2 months, and 2-6 months after the invite date than control.*

4 RESULTS

As Table 1 shows, Teahouse invitees were significantly more likely to survive in all 3-4 weeks after registration ($\chi^2 = 3.9$, $df = 1$, $p = 0.0479$, two-tailed) and to have made 5 or more edits between 2-6 months after registration ($\chi^2 = 4.6468$, $df = 1$, $p = 0.0311$, two-tailed). Editors in the experimental group were more than 10% more likely than those in the control to have met the minimum survival threshold in each of these survival windows.

Results from tests of number of edits (1+) between 1-2 months and 2-6 months, and number of edits (5+) between 1-2 months were insignificant, but they pointed in the same direction. A larger proportion of editors in the experimental group met the minimum survival criteria than those in the control group across all measurement periods. Although these results were not significant at

	1+ edits		
	3-4 weeks	1-2 months	2-6 months
control	247 (8.0%)	360 (11.6%)	397 (12.8%)
invited	1068 (9.1%)	1512 (13.0%)	1651 (14.1%)

	5+ edits		
	3-4 weeks	1-2 months	2-6 months
control	117 (3.8%)	191 (6.2%)	229 (7.4%)
invited	496 (4.2%)	841 (7.2%)	1008 (8.6%)

Table 1: Surviving new editors with the surviving proportion (XX%) are presented for each of the trial/survival periods and for both the 1 edit and 5 edit thresholds. Significantly different proportions are bolded.

$p < 0.05$, their directionality and p values suggest that the significant results presented above are not a statistical fluke.

5 DISCUSSION

The results of our study demonstrate that intervening early in a new Wikipedia editor’s career by presenting them with opportunities for socialization and social support from more experienced editors can increase their likelihood of contributing to Wikipedia over the subsequent weeks and months.

Importantly, the Teahouse is the only intervention of this kind that has been shown to be effective at increasing new editor retention on English Wikipedia in a controlled study, despite an extensive experimentation with a wide variety of innovative social and technical solutions developed by both Wikipedia community members and the Wikimedia Foundation.

We posit several possible reasons why the Teahouse was effective where other interventions have failed. First, the Teahouse provides support that is tailored to the new editor’s individual needs. Like Adopt-a-User, the Teahouse pairs newcomers with experienced and friendly Wikipedia community members who understand Wikipedia’s complexity and the challenges that this complexity presents to new editors. Unlike the Wikimedia Foundation Growth team’s experimental UI interventions, which were designed to address specific, hypothesized needs of brand new editors, the Teahouse is capable of addressing many different types of new editor challenges on an ongoing, as-needed basis—even providing them with the opportunity to become Teahouse hosts and support other newcomers once they had established themselves within the community. However, like the Wikimedia Foundation Growth Team interventions, the Teahouse is a scalable solution which can provide support to a large percentage of the new editors who join Wikipedia every day.

5.1 Limitations

In this study, we analyzed the effect of invitation to the Teahouse, rather than participation in the Teahouse. This was both an ethical and a pragmatic decision on our part. Wikipedia is a living system filled with real people with real problems and important work to do, and governance of Wikipedia is shared between the editing community and the Wikimedia Foundation. A true control

for the impact of Teahouse participation, as opposed to invitation, would have required either the development and maintenance of a non-Teahouse alternative designed to provide an intentionally deficient experience for new editors, or the introduction of deception within the study design. Either option would have resulted in more discouraged newcomers and involved wasting the time of volunteer Wikipedia contributors.

The study design used here is similar to an ‘intention to treat’ analysis, commonly employed in clinical trial settings where it is not possible to ensure that all members of the experimental group will comply with the study protocol [13]. As such, our analysis is inherently conservative, and so is our interpretation of our results.

As only 331 of the editors in the experimental group subsequently participated in the Teahouse during the study period, it is likely that the effect we observed is at least partially attributable to the invitation alone. Unlike the new editors in the control group, these newcomers received a personalized welcome message signed by a real Wikipedia editor, in addition to any of the generic, confusing, and often intimidating messages that they may have received from other Wikipedia editors or bots during this period.

We believe this is actually a validation of the Teahouse approach to new editor socialization. Considering the overwhelmingly negative and alienating experience that most new editors have when they first join Wikipedia, knowing that you are explicitly invited to participate, and that you have the opportunity to learn positive and useful things about how Wikipedia works and the community behind it, can be powerful motivators to continue contributing in themselves.

Being personally welcomed into a community—in effect, being labeled a member of the ‘in’ group—may encourage identity-based commitment to the group, which can increase retention even if the new member has not yet formed strong bonds with other community members [24]. Furthermore, finding help and mentorship opportunities is a perennial struggle for newcomers in open collaborations [27], and socialization is an ongoing process that extends beyond the initial orientation phase [29]. Making newcomers aware that support opportunities will be available to them when needed may increase their confidence and sense of self-efficacy [15].

New editors who received a Teahouse invitation were also given the opportunity to ‘lurk’ on the Teahouse without participating. Lurking has been shown to be a common and effective mechanism for legitimate peripheral participation in online communities [23]. Invitees who visited the Teahouse and browsed through recent questions profiles would have been able to learn about the background, interests, and experiences of other recent newcomers and experienced editors, witness many examples of positive and productive interactions between new editors and veterans, and potentially found answers to nascent questions they hadn’t yet thought to ask—all without making an edit.

In future work, it may be possible to perform a comparative analysis of the experience of active participation in Teahouse-style socialization spaces by taking advantage of pseudo-control conditions that occur naturally in Wikipedia and other open collaborations. On English Wikipedia, such an experiment might involve routing a control group to the Wikipedia Help Desk, a general purpose Q&A forum that is not specifically designed for new editors and which

lacks many of the Teahouse's social and technical affordances for newcomer socialization.

5.2 Next steps

Although most research on new editor experience of Wikipedia, and most Wikipedia research in general, has been performed on English Wikipedia, many of the other Wikipedia language editions have also seen a decline in their active editor populations in recent years. We are not aware of any research that has investigated whether the decline in these Wikipedias is also mediated by a decrease in new editor retention. However, recent interview and survey-based research performed by the Wikimedia Foundation on two mid-sized Wikipedias (Czech and Korean) shows that new editors on these projects face many of the same sociotechnical challenges as those faced by English Wikipedia newcomers [8]. This suggests that interventions like the Teahouse could be effective mechanisms for increasing new editor retention on these projects as well.

The majority of editor actions and interactions on Wikipedia result in an edit, and contribution patterns among new Wikipedia editors exhibit strong regularities across time. We believe that the new editor survival metric we used in this study provides a robust, if conservative, measure of absolute retention across different time scales, and recommend its use alongside other appropriate metrics to facilitate meaningful comparison of socialization-focused interventions.

Additional analysis of the factors that mediate the effect of Teahouse invitation or participation on retention—such as whether the new editor has their edits reverted, receives warning messages, or makes particular types of edits—before or after invitation, as well as interviews, surveys, and textual analysis of Teahouse questions and answers could help us understand what kind of socialization challenges the Teahouse is most effective at addressing, and suggest opportunities for additional interventions focused on particular issues that impact retention.

Finally, the effectiveness of the Teahouse and other similar social support forums could be improved by earlier intervention and better targeting. Currently, editors are only invited to the Teahouse after they have made at least 5 edits, and have been registered for at least 24 hours. The delay and edit requirement allow a degree of community vetting to take place before the decision to invite is made: new editors who exhibit malicious or disruptive behavior are identified by the community and filtered out. This reduces the chances of disruption on the Teahouse itself, and avoids the impression that bad actors are being rewarded for their disruptive behavior.

However, given the steep drop-off in retention first identified by Panciera et al [21], both the 24 hour delay and the 5-edit threshold likely result in many good-faith newcomers being denied the opportunity for positive socialization. If Wikipedia had a better way of accurately and efficiently discriminating between good-faith newcomers and vandals after less time and fewer edits had elapsed, a forum like the Teahouse could provide more new editors with timely and contextually relevant socialization.

A machine learning platform recently developed by the Wikimedia Foundation called ORES (Objective Revision Evaluation Service)⁵ provides powerful predictive models that can accurately discriminate between damaging and a non-damaging edits, and between malicious edits and edits that were made in good faith, even if they introduce errors or fail to comply with policies.

These models could be used to identify promising new editors who are struggling to learn the ropes of Wikipedia and invite them to the Teahouse in nearly real-time, reducing the likelihood that these editors will abandon the community before learning that support opportunities exist.

6 CONCLUSION

In this study we presented an evaluation of a forum on Wikipedia that is capable of providing a large number of new Wikipedia editors with positive and personalized support at multiple points in their socialization process. We performed a controlled study of the impact of Teahouse invitations on both short-term and long-term new editor retention and found that Teahouse invitations increased retention for both low- and high-activity newcomers in the short term (3-4 weeks), and for high-activity newcomers in the long term (2-6 months). We introduced a robust, flexible metric of new contributor survival that can be used in other investigations of Wikipedia new editor activity to facilitate cross-comparison and benchmarking. We presented a discussion of considerations for the design of responsible and effective interventions (and evaluations) of living open collaboration systems, and highlighted several opportunities for building from and validating the work presented here.

REFERENCES

- [1] Mark S. Ackerman and David W. McDonald. 1996. Answer Garden 2: Merging Organizational Memory with Collaborative Help. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (CSCW '96)*. ACM, New York, NY, USA, 97–105. <https://doi.org/10.1145/240080.240203>
- [2] Ivan Beschastnikh, Travis Kriplean, and David W McDonald. 2008. Wikipedian Self-Governance in Action: Motivating the Policy Lens. In *ICWSM*.
- [3] Moira Burke, Robert Kraut, and Elisabeth Joyce. 2010. Membership Claims and Requests: Conversation-Level Newcomer Socialization Strategies in Online Groups. *Small Group Research* 41, 1 (2010), 4–40. <https://doi.org/10.1177/1046496409351936> arXiv:<https://doi.org/10.1177/1046496409351936>
- [4] Boreum Choi, Kira Alexander, Robert E. Kraut, and John M. Levine. 2010. Socialization Tactics in Wikipedia and Their Effects. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW '10)*. ACM, New York, NY, USA, 107–116. <https://doi.org/10.1145/1718918.1718940>
- [5] Giovanni Luca Ciampaglia and Dario Taraborelli. 2015. MoodBar: Increasing New User Retention in Wikipedia Through Lightweight Socialization. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & #38; Social Computing (CSCW '15)*. ACM, New York, NY, USA, 734–742. <https://doi.org/10.1145/2675133.2675181>
- [6] Nicolas Ducheneaut. 2005. Socialization in an open source software community: A socio-technical analysis. *Computer Supported Cooperative Work (CSCW)* 14, 4 (2005), 323–368.
- [7] Andrea Forte and Cliff Lampe. 2013. Defining, Understanding, and Supporting Open Collaboration: Lessons From the Literature. *American Behavioral Scientist* 57, 5 (2013), 535–547. <https://doi.org/10.1177/0002764212469362> arXiv:<https://doi.org/10.1177/0002764212469362>
- [8] Lauren Gardener, Panthea Lee, Abigail Ripstra, and Neil Patel Quinn. 2017. New Editor Experiences. https://www.mediawiki.org/wiki/New_Editor_Experiences
- [9] R Stuart Geiger, Aaron Halfaker, Maryana Pinchuk, and Steven Walling. 2012. Defense Mechanism or Socialization Tactic? Improving. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*.

⁵<https://www.mediawiki.org/wiki/ORES>

- [10] Aaron Halfaker. 2013. Wikimedia Foundation Growth Team experiments: Onboarding new Wikipedians. https://meta.wikimedia.org/wiki/Research:Onboarding_new_Wikipedians
- [11] Aaron Halfaker, R. Stuart Geiger, Jonathan T. Morgan, and John Riedl. 2013. The Rise and Decline of an Open Collaboration System: How Wikipedi-afis Reaction to Popularity Is Causing Its Decline. *American Behavioral Scientist* 57, 5 (2013), 664–688. <https://doi.org/10.1177/0002764212469365> arXiv:<https://doi.org/10.1177/0002764212469365>
- [12] Aaron Halfaker, R. Stuart Geiger, and Loren G. Terveen. 2014. Snuggle: Designing for Efficient Socialization and Ideological Critique. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 311–320. <https://doi.org/10.1145/2556288.2557313>
- [13] John M Lachin. 2000. Statistical considerations in the intent-to-treat principle. *Contemporary Clinical Trials* 21, 3 (2000), 167–189.
- [14] Cliff Lampe and Erik Johnston. 2005. Follow the (Slash) Dot: Effects of Feedback on New Members in an Online Community. In *Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work (GROUP '05)*. ACM, New York, NY, USA, 11–20. <https://doi.org/10.1145/1099203.1099206>
- [15] Cliff Lampe, Rick Wash, Alcides Velasquez, and Elif Ozkaya. 2010. Motivations to Participate in Online Communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1927–1936. <https://doi.org/10.1145/1753326.1753616>
- [16] Jonathan T Morgan. 2012. Teahouse phase two report. https://meta.wikimedia.org/wiki/Research:Teahouse/Phase_2_report
- [17] Jonathan T. Morgan, Siko Bouterse, Heather Walls, and Sarah Stierch. 2013. Tea and Sympathy: Crafting Positive New User Experiences on Wikipedia. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*. ACM, New York, NY, USA, 839–848. <https://doi.org/10.1145/2441776.2441871>
- [18] Jonathan T. Morgan and Mark Zachry. 2010. Negotiating with Angry Mastodons: The Wikipedia Policy Environment As Genre Ecology. In *Proceedings of the 16th ACM International Conference on Supporting Group Work (GROUP '10)*. ACM, New York, NY, USA, 165–168. <https://doi.org/10.1145/1880071.1880098>
- [19] David R. Musicant, Yuqing Ren, James A. Johnson, and John Riedl. 2011. Mentoring in Wikipedia: A Clash of Cultures. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym '11)*. ACM, New York, NY, USA, 173–182. <https://doi.org/10.1145/2038558.2038586>
- [20] Sneha Narayan, Jake Orlovitz, Jonathan Morgan, Benjamin Mako Hill, and Aaron Shaw. 2017. The Wikipedia Adventure: Field Evaluation of an Interactive Tutorial for New Users. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, New York, NY, USA, 1785–1799. <https://doi.org/10.1145/2998181.2998307>
- [21] Katherine Panciera, Aaron Halfaker, and Loren Terveen. 2009. Wikipedians Are Born, Not Made: A Study of Power Editors on Wikipedia. In *Proceedings of the ACM 2009 International Conference on Supporting Group Work (GROUP '09)*. ACM, New York, NY, USA, 51–60. <https://doi.org/10.1145/1531674.1531682>
- [22] Maryana Pinchuk. 2012. Wikimedia Foundation Growth Team experiments: Post-edit feedback. https://meta.wikimedia.org/wiki/Research:Post-edit_feedback
- [23] Jenny Preece, Blair Nonnecke, and Dorine Andrews. 2004. The top five reasons for lurking: improving community experiences for everyone. *Computers in human behavior* 20, 2 (2004), 201–223.
- [24] Yuqing Ren, Robert Kraut, Sara Kiesler, and Paul Resnick. 2012. Encouraging commitment in online communities. *Building successful online communities: Evidence-based social design* (2012), 77–124.
- [25] Wikimedia Foundation Research. 2014. Research Metric standardization: Surviving new editor. https://meta.wikimedia.org/wiki/Research:Surviving_new_editor
- [26] Ellen L.Rubenstein. 2015. Rituals of introduction and revolving roles: socialization in an online breast cancer community. *Library & Information Science Research* 37, 4 (2015), 353–362.
- [27] Igor Steinmacher, Tayana Conte, Marco Aurélio Gerosa, and David Redmiles. 2015. Social Barriers Faced by Newcomers Placing Their First Contribution in Open Source Software Projects. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & #38; Social Computing (CSCW '15)*. ACM, New York, NY, USA, 1379–1392. <https://doi.org/10.1145/2675133.2675215>
- [28] I. Steinmacher, I. Wiese, A. P. Chaves, and M. A. Gerosa. 2013. Why do newcomers abandon open source software projects?. In *2013 6th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE)*. 25–32. <https://doi.org/10.1109/CHASE.2013.6614728>
- [29] Igor Steinmacher, Igor Scaliante Wiese, Tayana Conte, Marco Aurélio Gerosa, and David Redmiles. 2014. The Hard Life of Open Source Software Project Newcomers. In *Proceedings of the 7th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE 2014)*. ACM, New York, NY, USA, 72–78. <https://doi.org/10.1145/2593702.2593704>
- [30] Bongwon Suh, Gregorio Convertino, Ed H. Chi, and Peter Piroli. 2009. The Singularity is Not Near: Slowing Growth of Wikipedia. In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration (WikiSym '09)*. ACM, New York, NY, USA, Article 8, 10 pages. <https://doi.org/10.1145/1641309.1641322>
- [31] Georg Von Krogh, Sebastian Spaeth, and Karim R Lakhani. 2003. Community, joining, and specialization in open source software innovation: a case study. *Research Policy* 32, 7 (2003), 1217–1241.
- [32] Steven Walling and Dario Taraborelli. 2012. Wikimedia Foundation Growth Team experiments: Account creation UX. https://meta.wikimedia.org/wiki/Research:Account_creation_UX
- [33] Dennis M. Wilkinson. 2008. Strong Regularities in Online Peer Production. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC '08)*. ACM, New York, NY, USA, 302–309. <https://doi.org/10.1145/1386790.1386837>
- [34] Haiyi Zhu, Amy Zhang, Jiping He, Robert E. Kraut, and Aniket Kittur. 2013. Effects of Peer Feedback on Contribution: A Field Experiment in Wikipedia. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2253–2262. <https://doi.org/10.1145/2470654.2481311>