

Wikipedians Are Born, Not Made

A Study of Power Editors on Wikipedia

Katherine Panciera, Aaron Halfaker, Loren Terveen
GroupLens Research
Department of Computer Science and Engineering
University of Minnesota
Minneapolis, Minnesota
{katpa,halfak,terveen}@cs.umn.edu

ABSTRACT

Open content web sites depend on users to produce information of value. Wikipedia is the largest and most well-known such site. Previous work has shown that a small fraction of editors – *Wikipedians* – do most of the work and produce most of the value. Other work has offered conjectures about how Wikipedians differ from other editors and how Wikipedians change over time. We quantify and test these conjectures. Our key findings include: Wikipedians’ edits last longer; Wikipedians invoke community norms more often to justify their edits; on many dimensions of activity, Wikipedians start intensely, tail off a little, then maintain a relatively high level of activity over the course of their career. Finally, we show that the amount of work done by Wikipedians and non-Wikipedians differs significantly from their very first day. Our results suggest a design opportunity: customizing the initial user experience to improve retention and channel new users’ intense energy.

Categories and Subject Descriptors

H.5.3 [Group and Organization Interfaces]: Computer-supported cooperative work, web-based interaction

General Terms

Human Factors

Keywords

Wikipedia, wiki, power editors, contribution, collaboration

1. INTRODUCTION

In open content online communities, users produce and share information of value. To succeed, these communities require active and committed members. However, research has shown that there is a wide spectrum of activity in these

communities. Some communities don’t attract enough attention to succeed. Even in those that do, there are typically a small number of dedicated participants and a much larger number of “lurkers” who do not participate [13]. These patterns have been observed in diverse forms of online communities implemented in a variety of technologies, including social networks, forums, wikis, Usenet groups, and mailing lists [10, 14, 17, 23, 35].

The English language Wikipedia is a large and very active open content community, with over 1.5 million registered users and over 2.7 million articles. The issue of who does the work in Wikipedia – the few or the masses – has been a subject of debate in blog conversations and research papers [2, 30, 33, 35]. By this time, it has become established that a small proportion of editors account for most of the work done (for example, see [19, 21]) and the value produced [28]. However, many issues about the work of these power editors, or *Wikipedians*, remain unknown.

Bryant took a qualitative look at Wikipedians [4]. They conducted interviews with nine active Wikipedia editors. They concluded that Wikipedians filled different niches than non-Wikipedia editors, branched out to new areas and topics as they matured, and gradually took on more “community-oriented” work such as enforcing policy or patrolling for vandalism. One of the contributions of our work is a quantitative test of Bryant’s conclusions.

Specifically, we investigate the quantity, quality, and nature of Wikipedian’s activity. We look at how their work changes over their lifespans in Wikipedia, and how it compares to less-prolific editors. One pattern emerged from many of our analyses: Wikipedians look different from other editors from the very beginning. Furthermore, in most respects Wikipedians do not increase or improve their activity over time. Thus, our title: Wikipedians are born, not made.

In the remainder of the paper, we survey related work, describe the data we analyzed and how we formalized the notion of a Wikipedian, present and discuss our results, then close with ideas for future work and a brief summary.

2. RELATED WORK

While we focus our research on Wikipedia, other open content online communities have been embraced and studied by researchers over the past two decades, helping us form many of our basic understandings of online content contribution systems.

2.1 Research In Other Online Communities

Our work deals with the notions of contributions of quantity and quality in open content communities, as well as the notion of community diversity and norms. These are ideas that have been studied in various online communities in the past.

Quantity: It has been widely observed that in online communities the minority of participants provide a majority of the content. This production of content often resembles a power law. This has been seen in Usenet postings [14, 34] as well in the blogging community [26].

Quality: Researchers have investigated quality in venues from Slashdot, where quality is determined through moderation [24], to the ratings of movies in MovieLens [7], but quality depends very much on context. Unlike these settings, much of Wikipedia’s content is factual, but quality is ultimately determined by editors, who, upon finding a statement that is incorrect or disagreeable to them, can simply delete or alter it.

Community: Diversity within online communities and the community norms surrounding the community can greatly influence the member experience, especially for outsiders. Dugan studied social network profiles and found results which suggest that a diverse profile positively influences the number of “friends” the user has on the site [12]. Ducheneaut investigated community norms in the Python community and learned these norms affect the experience of the newcomer and how the process of joining the community requires successful completion of several rites of passage [11].

2.2 Wikipedia

Wikipedia has been of interest to researchers since it was first introduced in 2001.

Many researchers have also investigated the role of conflict in Wikipedia. Viégas led the way with her seminal paper about vandalism and visualizations on Wikipedia [31]. Kriplean found that citing policy in arguments does not lead to the resolution of all conflicts [22] and Suh and Kittur both worked on developing visual models for Wikipedia conflict [29, 21].

2.2.1 Quantity of Work

Researchers interested in the quantity of work editors generate on Wikipedia have most often looked at quantity over the life of Wikipedia. Kittur, Almeida, and Ortega all found that contributions have increased dramatically [19, 20, 2, 27].

These other researchers have all looked at contributions over the life of Wikipedia, while we look at contributions over the life of an editor. Since the user base of Wikipedia has grown immensely since the early days, it makes sense that the quantity of edits has also increased over that time period. By analyzing edits over the life of an editor, we remove the effects of the size of the user base and the entry date of the editor.

2.2.2 Quality of Work

Few researchers have investigated the notion of quality on Wikipedia. In 2005, Giles, in a Nature article, found that the accuracy of science articles on Wikipedia was very close to the accuracy of the same articles in the Encyclopedia Britannica by employing human judges [15].

Others have used word persistence as a proxy for content

quality. Adler and Alfaro developed a reputation metric for editors which measured how long an editor’s changes last over time[1]. Priedhorsky found that the top 10% of editors (by number of edits) contribute 86% of the value when measured by word views on the English language Wikipedia and that an even more elite group, the top 0.1% by number of edits (about 4400 editors) contribute 44% of the value [28].

Like Priedhorsky and Adler and Alfaro we use persistence as a proxy for quality. While Priedhorsky looked at persistence from a reader’s perspective, we, like Adler and Alfaro, look at the persistence of words from an editor’s perspective. We approach the persistence of content slightly differently by interpreting subsequent revisions as an informal peer review. This allows us to measure the quality of a contribution by examining its acceptance or rejection by the rest of the Wikipedia community.

2.2.3 Community Work

The encyclopedia articles of Wikipedia are only one portion of the editable content. Wikipedia divides the content into nine distinct namespaces. They are divided based on the function that they serve within the system. For each of these content namespaces, there is an accompanying “Talk” namespace that is designated for communication and coordination.

In terms of community work, researchers have been particularly interested in Wikipedia Talk pages, specifically in policy usage and coordination among the Wikipedia community. Kittur found that in 2001, 90% of edits were done in the Main namespace on Wikipedia but that this number dropped to 70% by June 2006, supporting the idea that over time, Wikipedians have diversified their efforts to other namespaces [21]. Wilkinson found that articles with more discussion on their Talk page were generally ranked higher in quality according article ratings [36]. Viégas found that over half of Talk page comments are requests for coordination and 8% are policy invocations, which, she argues, leads to the general conclusion that Wikipedia contains strong and supportive communities [32]. Butler also found that much of the explicit coordination is managed through the Talk or discussion pages for the article in question [6].

Like Kittur, we are looking at the percentage of edits that are done in different namespaces on Wikipedia, but unlike Wilkinson, Viégas, and Butler, we do not analyze the discussion on the Talk pages.

3. DATA AND DEFINITIONS

3.1 Where Did We Get Our Data?

The data used in this paper was from the January 13, 2008 complete English language Wikipedia dump for all namespaces.¹ Edits by known *bots* (autonomous software programs) were removed, as were edits made by anonymous editors or with the AutoWikiBrowser (AWB). AWB is a semi-automated editor that helps Wikipedia editors complete a series of repetitive tasks, such as formatting, more quickly and easily. We excluded *bots* and users with AWB access because they make abnormal numbers of edits due to their autonomous and semi-automated natures, respectively. For example, AntiVandalBot, a *bot* that cleans up

¹<http://download.wikimedia.org/enwiki/20080113/enwiki-20080103-pages-meta-history.xml.7z>

vandalism, can make up to three edits in a minute, and edit for hours at a time, making over 500 edits in a 20 hour period. We feel that *bots* should be not be grouped with humans in analyses and chose to focus on Wikipedians and non-Wikipedians, not Wikipedians, non-Wikipedians, and *bots* and AWB users.. All of the data for this paper was parsed from the dump and entered into a database. The graphs and statistics were generated using R.

3.2 Who Did We Count?

The term *Wikipedian* is used informally within the Wikipedia community and ranges in definition from a registered editor to a small, select subset of editors. Using it within a research study requires a precise definition. Bryant characterized Wikipedians as editors that had been active on average 14 months and reported daily or near daily activity.

We define Wikipedians as editors who have made at least 250 edits over their lifetime. This is the threshold required for registering for counter-vandalism tools including VandalProof², and is half of the figure required to obtain AWB access. Our Wikipedians as a group edit 28% of the days between their first registered edit and their last edit, while our non-Wikipedia group was only active 4.3% of the days.

We did a number of things to make sure our definition was not harmfully arbitrary. First, we varied the threshold and tried our analyses to see if the patterns changed. In nearly all cases, the answer was ‘no’. However, in a few cases – for “super-elite” editors who made more than 5000 edits – it did. In these cases, we do present results using more thresholds. Second, we considered an alternative perspective for defining Wikipedians: length of activity rather than amount of activity. Thus, we looked at definitions like: editors who had at least one or two years between their first and last edits. However, when we did analysis with these “long term” editors, their activity was quite similar to non-Wikipedia editors. Thus, we retained the definition presented above.

We only considered registered users for this study. When registered users log in to Wikipedia, any edits they made are tagged with their user name. Edits of non-logged-in users are tagged by IP address. A single IP address may include edits from multiple people (if they all used the same machine, or, at one time, if they came to Wikipedia using the AOL service), and a single person might have multiple IP addresses (if he/she used different machines).

Our Wikipedian sample is made up of all 37,956 registered editors that fit our definition as of early January 2008. For comparison purposes, we want an equal sized set of *non*-Wikipedians. Thus, we also include a random sample of 38,975 editors from the remaining 1.5 million editors. Why don’t we have equal numbers of Wikipedians and non-Wikipedia editors? Initially, we did. However, we then realized that our Wikipedians contained some bots and AWB users, which we removed. We did not generate another random sample of non-Wikipedia editors due to the extensive computational time this required and the small benefit it offered.

3.3 What Did We Count?

Our basic unit is *edits per day per editor*. Each editor’s edits were grouped according to when his or her first edit occurred. For example, if an editor made his first edit at

²<http://en.wikipedia.org/w/index.php?title=User:AmiDaniel/VandalProof&oldid=85247378>

11:15am, we considered his second day of editing to start 24 hours later (at 11:15 am), his third day of editing to start 48 hours later (again at 11:15 am), etc.

Although in this paper, we refer to this time as the day in the editor’s life, it is possible that they could have made anonymous edits prior to registering or could have made edits with a prior account. In this case, their first edit as a registered user may not be their first edit to Wikipedia. Although we refer to this as the day in the editor’s life, it really represents the day in the life of this editor’s account. Our work does not rely on any assumption of when the editor actually began editing Wikipedia. This does not invalidate our results as we are still able to make predictions from their first day as a registered editor.

We considered an alternative to the edit, namely the *session*. The intuition here is that editors might make multiple edits to the same article in quick succession. We formalized the notion of a session as: multiple consecutive edits to a single article, with no edits to other articles and no edits by others to the said article, within one hour. We found that 75% of users had an average of 1.4 edits per session or fewer. Many of these repetitive edits occur when editors use the save button to preview their edit or realize after they’ve saved something that they’ve made a mistake. Because most sessions amount to just about a single edit, and edits are simpler and more tangible, we decided to stay with the edit as our unit.

4. RESULTS AND DISCUSSION

4.1 How Results Are Presented

Most of our results are presented in graphs, many of which share a number of key characteristics. To aid understanding, we describe these characteristics here. All graphs have days (of editors’ lifespans) on the X axis. The X axis uses a \log_2 scale, making it easier to see what is happening in the first few days of editors’ lifespans. Error bars are shown for points; however, in most cases, the errors are so small that the bars are not visible.

The graphs were produced in color and are easier to understand if viewed in color. The captions refer to colors where relevant. However, we also tried to make the graphs comprehensible if viewed in black and white.

Most graphs show Wikipedians (black) and non-Wikipedians (red). Several graphs use a more detailed breakdown of editors. In these cases, the caption and surrounding text specify the breakdown.

Finally, the daily average for all analyses will include many “zero editors”, i.e., editors who did not make any edits on that day. We continued counting an editor in computing the average for a day until the day of their last edit, at which time we no longer included them. Figure 1 shows the percentage of Wikipedians and non-Wikipedians who make an edit on any given day.

4.2 Wikipedians and Work Quantity

By our definition, Wikipedians do more work than non-Wikipedia editors overall. However, we want to examine work quantity in more detail. In particular, we want to look at (1) how the amount of work done changes (or stays constant) over an editor’s “lifespan” in Wikipedia, (2) whether the difference in work quantity between Wikipedians and non-Wikipedians changes or stays constant over time, and

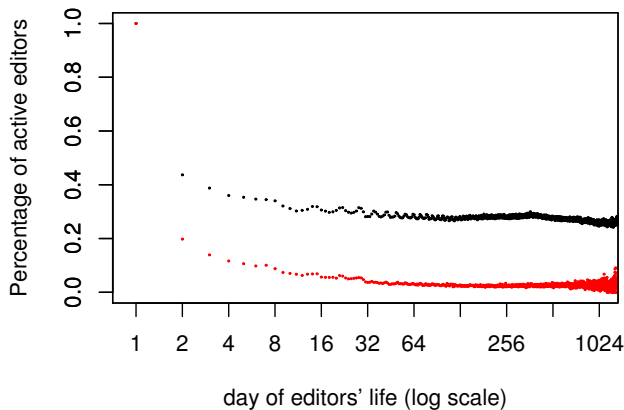


Figure 1: Percentage of users who made an edit on the given day. The top black scatter shows the Wikipedians, while the bottom red scatter shows non-Wikipedian editors. Both populations start at 1. Standard error bars are displayed.

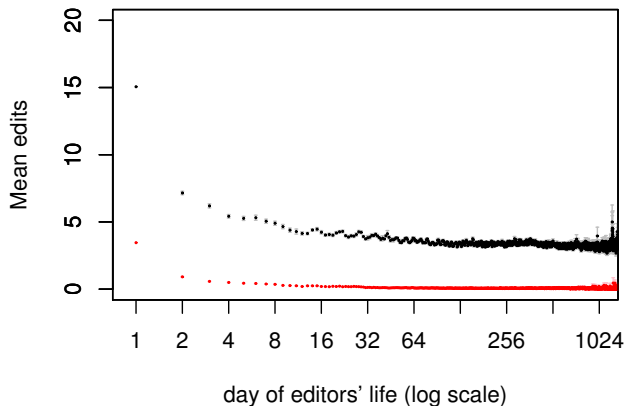


Figure 2: Average number of edits by day across all namespaces. The top black scatter represents Wikipedians and the bottom red scatter represents non-Wikipedian editors. Standard error bars are displayed.

(3) what the earliest days of editors' activity tells us about the prospects for their later activity.

Figure 2 shows edits per day for both Wikipedians and non-Wikipedian editors. From the graph we can make several observations. Both groups of editors begin with a burst, then tail off rather quickly to a level that they more or less maintain for the rest of their lifespans. However, Wikipedians do much more work than non-Wikipedian editors at every stage. Wikipedians made 15.1 edits in their first day, while non-Wikipedian editors made 3.5 edits in their first day. After two months, Wikipedians are down to about 3 edits a day, while non-Wikipedians essentially do nothing. This is a rather dramatic discontinuity: when you try Wikipedia, you either do a lot of work over time, or do a little and quickly lose interest.

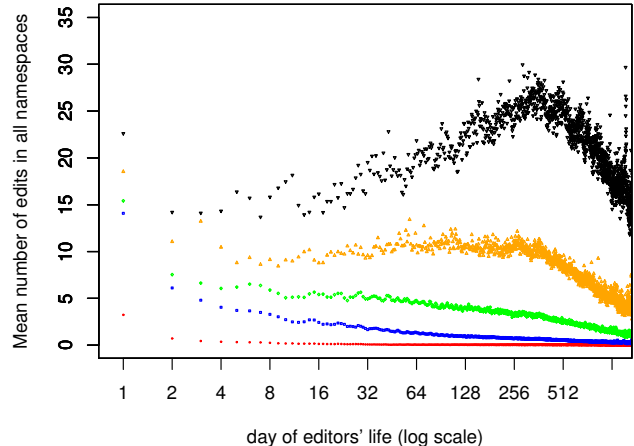


Figure 3: Number of edits by day across all namespaces. The top black scatter represents editors with over 10000 edits, the next down orange scatter is editors with 5001-10000 edits, green scatter in the middle is editors with 1001-5000 edits, blue scatter near the bottom is editors with 101-1000 edits, and the red scatter at the bottom is editors with 1-100 edits.

Two other analyses elaborate this insight. First, in Figure 3 we refined our definition of Wikipedian to consider more and finer grained levels of activity. Previous research [19, 27] segmented editors into the following groups: less than 100 edits, 101 to 1000 edits, 1001 to 5000 edits, 5001 to 10,000 edits, and over 10,000 edits. Figure 3 shows the average number of edits per day with these groups of editor. All the groups up to 5000 edits had the same pattern seen in Figure 2: initial burst followed by decline.

However, the most active editors (about 3600 in the top two groups) had a different pattern. Editors with 5001-10,000 edits stayed roughly constant for about the first year, and editors with over 10,000 edits **increased** their activity from about the first month well into the second year of their lifespans. The extreme right end of the graph is noteworthy in two respects: (1) error bars become noticeable because the number of active editors becomes quite low, and (2) on average, even the most prolific editors do fade out given enough time.

A second analysis illustrates how dramatically the activity of Wikipedians and non-Wikipedian editors differs from the beginning. Table 1 shows that the number of edits made in the first two days is a strong predictor of the probability that a new editor will become a Wikipedian. Looking at the first day, note that less than 1% of those making a single edit eventually become Wikipedians, 4.5% of those making 6-10 edits do, and over 8% of those making 11-20 edits do. However, the second day's activity is even more informative. Making just a single edit yields nearly a 6% probability of becoming a Wikipedian, and making 6-10 edits means over 18% probability.

4.2.1 Discussion

Nearly all editors begin with a burst of activity, then quickly tail off. Non-Wikipedians tail off to almost no edits

Number of Edits	First Day		Second Day	
	Proportion of Editors	Likelihood of being Wikipedian	Proportion of Editors	Likelihood of being Wikipedian
0 (no subsequent edits)	NA	NA	62.21%	.00021%
0	NA	NA	29.74%	4.47%
1	45.57%	.95%	2.84%	5.62%
2-3	29.11%	1.62%	2.21%	8.83%
4-5	10.25%	2.72%	.93%	12.83%
6-10	8.67%	4.50%	.98%	18.19%
11-20	4.16%	8.27%	.61%	26.75%
21-40	1.59%	15.42%	.33%	37.29%
Over 40	.65%	31.11%	.15%	64.82%

Table 1: Likelihood, based on first two days of edits, of a user making a given number of edits and becoming a Wikipedian. Results are not cumulative and are independent.

by day 16, which suggests that an exceptional amount of work is done by editors in their first few weeks.

Overall, 60% of registered users never make another edit after their first 24 hours. There are several possible reasons for this. First, editors might be scared away by negative reactions to their edits such as outright removal of their added content. Second, editors might not be engaged by the existing community. Both of these reasons are negative for the community. Another reason for users to leave is that they may have registered in order to complete a single one-time task. This is not a negative reflection on the community.

The contribution rates of Wikipedians drop off over time. They do not fall to zero, as the non-Wikipedians do, but rather hover around four. Unlike the non-Wikipedians, they tend to remain more active. If we can figure out why they stick around, maybe we can learn more about how to retain non-Wikipedians.

These results suggests several design ideas and possibilities for future work.

First and most obvious, try to get new editors to come back! This could be done by automatically selecting the feature to have email notifications sent when someone posts on an editor’s talk page or automatically checking the box to add edited pages to watchlist and sending email notifications of updates to watched pages. This is a common practice in communities such as LinkedIn and Facebook. Prior work shows that carefully crafted messages can be quite successful in getting users to return and do work in an online community [3, 16].

Second, look for ways to direct the burst: specifically, point new editors to work that they tend to be most successful with, whether filling in desired content, checking facts in disputed articles, or learning about Wikipedia conventions and editing articles to conform to them. In addition to getting more work done, this has an additional important benefit: by routing users into different niches, it’s possible for them to feel they are making a substantial contribution. And a robust finding in social psychology [18] and social computing [3, 25] is that making users believe they can make a significant contribution increases the amount of work they do and can improve their retention. Cosley found that welcome messages help with sustaining an engaging community and increasing retention [8]. In addition, he implemented a recommender system to help people find work to do on Wikipedia and found that the customized to-do list elicited four times as many edits as a random list of articles [9].

Other online communities can benefit from the observation that power users can be noticed early on in their lives as shown in Table 1. This is important both for picking moderators and administrators and in not ostracizing those who will become powerful contributors later in their lives. Future research could look for the first and second day retention effects in other online communities.

4.3 Wikipedians and Work Quality

If most Wikipedians do not increase the **amount** of work they do over time, perhaps they increase the **quality**. We checked this next.

No single, universal quality metric exists for Wikipedia articles. There are a few candidates. Wikipedia articles can be rated by Wikipedia editors, and some prior work has used these ratings. Human coding can be used to judge quality [15] but is time intensive and cannot be extrapolated to a wider portion of Wikipedia, so in this work, we take another approach. Following [1] and [28], we used **persistence** as a proxy for quality. Intuitively, it is reasonable to assume that the more of an editor’s content lasts, and the longer it lasts, the higher its quality. To formalize this, we calculated the average number of revisions that the words added by all the editors in our sample lasted.

We chose our method in part to use an editor based metric and in part to be able to assign value per word added. The Wikipedia Assessment Ratings (as used by Kittur [19]) are internal to Wikipedia and are a reader-based metric that is displayed to users and based on a peer review process. The ratings are also assigned per article instead of per edit or per word. Priedhorsky’s PWV metric was also a reader-based metric measuring the number of views that a given word received [28]. Adler and Alfaro used an editor-based metric that assigned value to words based on how long they had lasted [1]. Our PWR metric is based on implicit approval of words added by subsequent editors and is an editor-based metric assigned per word.

Revision	Editor	Text
1	Steve	blue apples are yummy
2	Chris	apples are yummy
3	Paul	apples are certainly yummy
4	Robin	apples are certainly most yummy
5	Phil	apples are nutritious

Table 2: Example Revision History

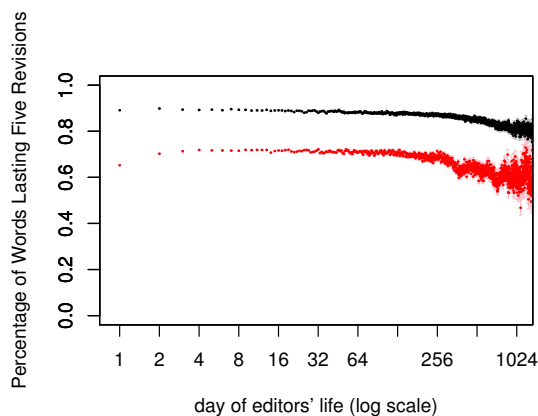


Figure 4: Wikipedians produce higher quality edits than non-Wikipedian editors. The top black scatter represents Wikipedians and the bottom red scatter represents non-Wikipedian editors. Standard error bars are displayed.

In the example in Table 2, Steve’s initial edit added four words. He will get one point for every revision that each word lasts. That is, he will get 0 revisions for blue, 4 for apples, 4 for are, and 3 for yummy. That gives him a net score for this edit of 11 persistent word revisions (*PWR*). Similarly, in this example, Chris has a score of 0, since she didn’t add any words, and Paul has a score of 1 since his (single) word lasts one revision.

We did analysis showing that if a word lasted for at least 4 revisions, 91% of the time it will last at least 10 revisions and 68% of the time it will last at least 50 revisions. Therefore, to compute our final score for an edit by a user, we simply counted the proportion of words that lasted at least 5 revisions.³ The “Main” namespace contains the encyclopaedia articles, so we only analyzed edits to this namespace.

Figure 4 shows that Wikipedians make higher quality edits than non-Wikipedian editors. This advantage is fairly large: in the steady state, Wikipedians average nearly 0.9 on our metric, and non-Wikipedian editor averages about 0.7. However, quality does not increase over time for either group; indeed, it actually decreases slightly.

We conducted the same analysis, showing mean PWR over lifetime for the editors, as bucketed by [19, 27]. The graph must be viewed in color. The graph shows that the PWR for users with under 100 edits is similar to the non-Wikipedians PWR, with the PWR for users with 101-1000 edits between that and the PWR for users with 1001-10000 edits. By day 64 of the editors’ lives, the editors with 1001-10000 edits have higher PWR (slightly) than the editors with over 10000 edits.

4.3.1 Discussion

These results are consistent with the quantity results of the previous section: Wikipedians seem to be born not made. They begin at a certain level of activity, well above that of non-Wikipedian editors, but they do not improve over time. However, we find the quality results more puzzling: Wikipedia editing is quite an odd activity, unlike, say, skiing or playing the piano or writing computer code or making

³For this analysis, we only considered revisions that would last at least 5 subsequent edits.

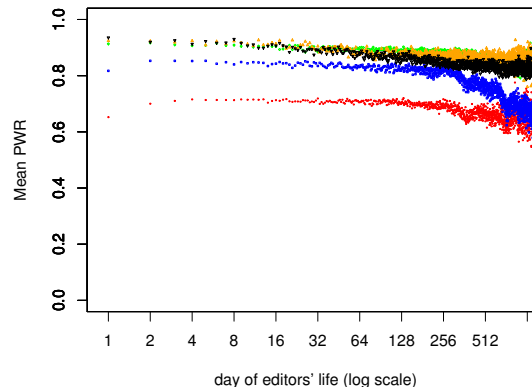


Figure 5: Wikipedians produce higher quality edits than average editor. The top black scatter represents editors with over 10000 edits, the next down orange scatter is editors with 5001-10000 edits, green scatter in the middle is editors with 1001-5000 edits, blue scatter near the bottom is editors with 101-1000 edits, and the red scatter at the bottom is editors with 1-100 edits. This graph is best viewed in color.

pancakes, if doing it 100s of times does not lead to any improvement. This raises several possibilities: Editors become more bold over time (making edits that are more controversial), or as Bryant suggests, editors edit items outside their expertise, or, although it seems unlikely, editors may get worse or lazy.

We might not be measuring actual quality, but we are measuring the perceived quality by other editors in Wikipedia. By measuring quality through other metrics, we may find different results. For example, if we use time based metrics, such as PWV or Adler and Alfaro’s persistence over time metric, we expect to see that there is a substantial advantage for early adopters. However if we use a metric such as Wikipedia Assessment Quality, we may see that as editors age, they become more active in higher rated articles.

Outside of Wikipedia, we know that experience can be a proxy for quality. For example, as seen in Ducheneaut, Python developers become better as the age [11]. This suggests that while some communities have a learning effect that influences quality, in others, quality remains static or unstable. Future research could investigate different online communities with quality measures to learn more about the effect of experience on quality.

4.4 Wikipedians and Community Work

One of the most intuitively appealing assertions in Bryant[4] is that Wikipedia editors began their careers by editing *content* on topics they knew about, but gradually shifted to doing more *community maintenance* editing. To test this assertion, we quantified it in several ways: (1) what *namespaces* were edited?, (2) did edits explicitly refer to community norms?

4.4.1 Namespace Diversity

As mentioned earlier, Wikipedia has 9 publicly editable content namespaces and 9 publicly editable communication namespaces, each one serving as the site for a different type

of activity. For our purposes, three namespaces are of interest – Talk, User Talk and Wikipedia – as areas for community maintenance activity. Talk is the namespace for discussing the content of Main articles; there is one Talk article for each Main page. Questions and conflicts about facts, requests for input, and any other conversation about a content article happens here. User Talk is for conversation about editors; each editor can have a User Talk page. User Talk pages are edited to welcome new users, to warn users that an edit was deleted for being vandalism, to discuss edits made by the user, and to acknowledge contributions by the user. The Wikipedia namespace is used to form and enforce policies, request and vote on admin-ship (a more powerful status an editor may attain), vote on banning editors, etc. 66% of all edits are made in Main, 9% in Talk, 8.4% in User Talk, and 6.5% in Wikipedia. These are the top four namespaces, and together they account for nearly 90% of all Wikipedia edits.

Bryant’s hypothesis would be supported if Wikipedians shifted their namespace “profile” over time; in particular, we expected they would edit less in Main and more in the other three namespaces. We also expected Wikipedians to be more active in these namespaces than non-Wikipedia editors; we expected non-Wikipedia editors to do essentially no community maintenance work.

However, the graphs of activity in the four namespaces – see Figures 6, 7, 8, and 9 – gave very little support to these hypotheses. Figure 6 shows that Wikipedians decreased their proportion of editing in Main over time only slightly. This slight decrease meant that only a slight increase in “community maintenance” namespaces was possible. And, indeed, Figures 7, 8, and 9 confirm this. The graphs for Wikipedia and User Talk tell similar stories. Wikipedians increased the proportion of activity devoted here very slightly, and they clearly devoted a larger proportion of their effort here than did non-Wikipedia editors.

The most puzzling pattern is seen for the Talk namespace. Since discussions of content (Main) articles occur here, this is the place where coordination is arranged and collaboration is managed [21, 32]. According to Wilkinson the more discussion on the Talk page, the higher quality the Main article will be [36]. However, our data (Figure 9) show that Wikipedians do not increase the proportion of work in the Talk namespace, nor do they do a higher proportion of their editing here than non-Wikipedia editors.

As in earlier analysis, we broke down the set of editors into finer-grained buckets: <100 edits, 101-1000 edits, 1001-5000 edits, 5001-10000 edits, and >10000 edits. Figure 10 shows that raw number of Talk edits per day by each group of editors. Again, the most active editors do increase the raw number of edits in Talk per day. However, a comparison with Figure 3 makes it clear that even they do not increase the proportion of effort devoted to Talk page.

4.4.2 Invocation of community norms

One often-noted aspect of Wikipedia is how the community of editors has evolved *norms* that govern its work. Probably the most famous of these is *Neutral Point of View* or *NPOV* which says that articles should be free of bias and written from a neutral perspective. Another important community maintenance activity is seeking out and reverting vandalism, i.e., malicious or mistaken edits that damage Wikipedia articles.

An edit has a comments field, which is a place where edi-

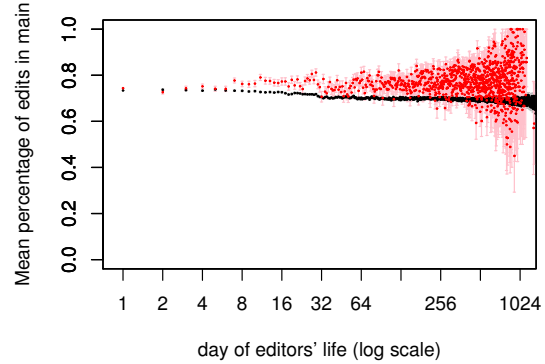


Figure 6: The percentage of edits in “Main” for all editors. The lower black scatter represents Wikipedians and the upper red scatter represents non-Wikipedia editors. Standard error bars are displayed.

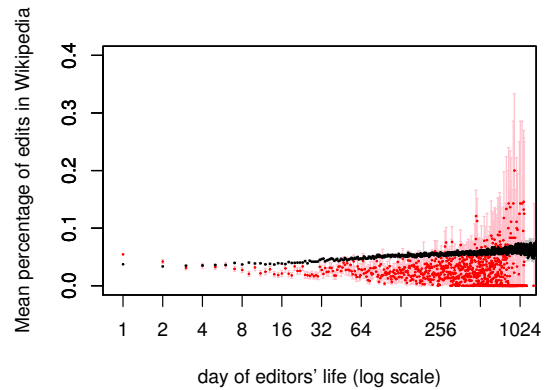


Figure 7: The percentage of edits in “Wikipedia” for all editors. The top black scatter represents Wikipedians and the bottom red scatter represents non-Wikipedia editors. Standard error bars are displayed.

tors may explicitly refer to norms. For example, “Revert per WP:NPOV” indicates a page has been reverted because it did not follow the *NPOV* policy. For our purposes, any comment text including “WP:” or “Wikipedia:”, which denotes a link to the namespace for policy and guidelines, was considered to be a policy comment. For detecting vandalism, we use the D-Loose method of Priedhorsky which detects 62% of vandalism.

Figure 11 shows that Wikipedians invoke norms more often than non-Wikipedia editors, Wikipedians become more likely to invoke norms, and that non-Wikipedia editors do not. This is the one place where we have seen a learning effect. Note that there are two possible types of learning that could be occurring here. First, Wikipedians could learn about and do more of the norm-enforcing activity (reverting vandalism or enforcing a policy). Second, perhaps they are doing the same amount of this activity, but are learning to note this in their comments. Either case indicates an enhanced understanding of community practices. Further research is necessary to distinguish the cases.

Figure 12 shows the same analysis for the finer-grained

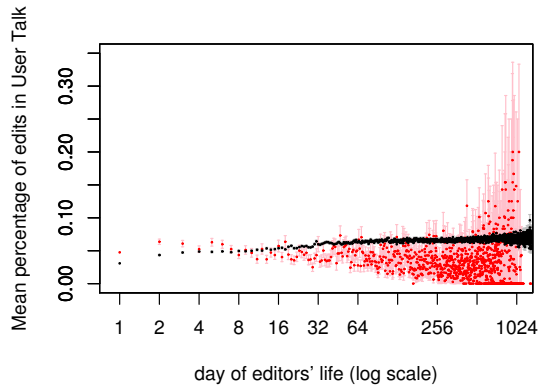


Figure 8: The percentage of edits in “User Talk” for all editors. The top black scatter represents Wikipedians and the bottom red scatter represents non-Wikipedian editors. Standard error bars are displayed.

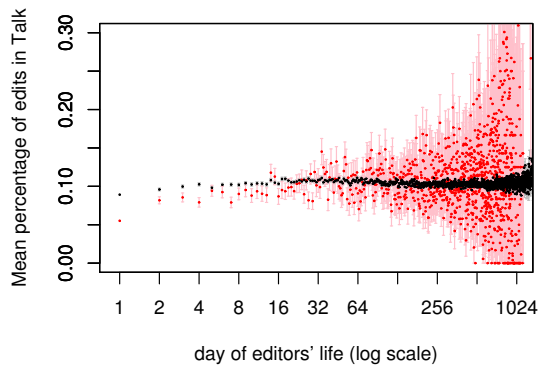


Figure 9: The percentage of edits in “Talk” for all editors. The top black scatter represents Wikipedians and the bottom red scatter represents non-Wikipedian editors. Standard error bars are displayed.

buckets. Although it is hard to tell in a black-and-white version of the figure, the learning effect goes all the way down to editors making between 101 and 1000 edits.

Another interesting thing to note about Figure 11 is that the average number of norm invocations on the first day is not 0. This indicates that at least some editors are aware of and making use of community norms from the very beginning of their Wikipedia careers.

4.4.3 Discussion

Of Talk pages and Wikipedians, Bryant writes “Although none of the interviewees described initial encounters with Wikipedia that involved discussion pages or page histories, these features became deeply integrated into their routine activities on the site” [4]. This seems intuitive and we do see that the Talk pages for Main are the most frequently used talk pages, however our results do not show a shift in activity towards Talk over an editor’s lifespan.

In our study we were not able to capture all communication. Users might move communications off of Wikipedia Talk pages after their initial activity, meaning that we are

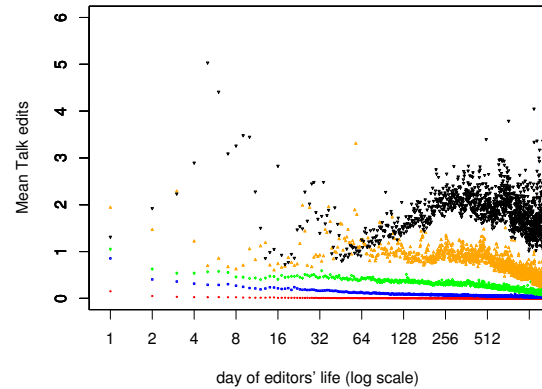


Figure 10: Average raw number of edits in “Talk” for all editors. The top black scatter represents editors with over 10000 edits, the next down orange scatter is editors with 5001-10000 edits, green scatter in the middle is editors with 1001-5000 edits, blue scatter near the bottom is editors with 101-1000 edits, and the red scatter at the bottom is editors with 1-100 edits. This graph is best viewed in color.

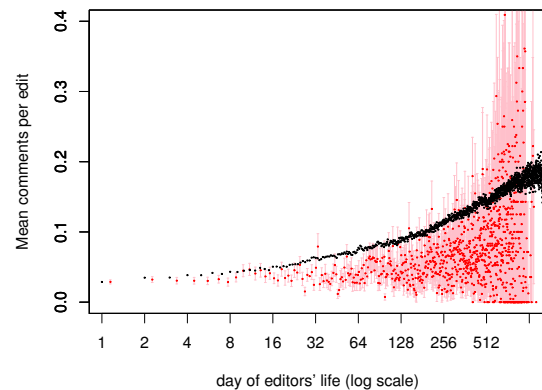


Figure 11: Average percentage of edits that include edit comments that reference Wikipedia policy or vandalism. The top black scatter represents Wikipedians and the bottom red scatter represents non-Wikipedian editors. Standard error bars are displayed.

not able to view their communications. Even so, it is highly unlikely that all of our subjects coordinate Wikipedia activity outside of the Talk pages.

If using Talk pages is important – and we think it is – then it might be useful for Wikipedia to “promote” Talk pages in its interface. For example, whenever someone opened the content editor for a Main page, some text could be added to the interface reminding them of the role of Talk page and presenting a link to click on the corresponding Talk page.

Online coordination is an issue for many communities. Think about ways to promote unused or underused features methods for communicating around topics this is obvious in forums etc

On the other hand, we found that Wikipedians did increase their invocation of community norms, supporting Bryant’s conclusion that “novice users learn the rules and conventions

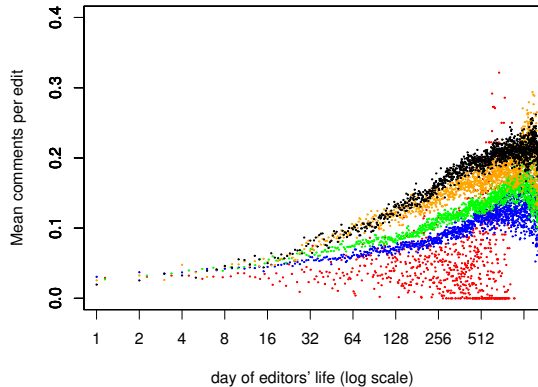


Figure 12: Average percentage of edits that include edit comments that reference Wikipedia policy or vandalism. The top black scatter represents editors with over 10000 edits, the next down orange scatter is editors with 5001-10000 edits, green scatter in the middle is editors with 1001-5000 edits, blue scatter near the bottom is editors with 101-1000 edits, and the red scatter at the bottom is editors with 1-100 edits. This graph is best viewed in color.

for contributing both through observation and direct coaching from more knowledgeable others” [4]. This could mean one of two things: either the editors are learning the community norms from scratch (going from not knowing about vandalism to citing vandalism as a problem) or editors are learning to invoke the related norms (going from reverting for NPOV to citing NPOV as the reason for the revert). Both of which are important for the community.

Burke found that the strongest predictors of a successful bid to become an administrator on Wikipedia included the number of edits in the Wikipedia namespace, as well as the number of edits in the Talk namespace and the number of edit comments that explicitly mention vandalism [5]. We found that the number of edits in the Wikipedia namespace and, in part, the number of edit comments referring to vandalism, are also different between Wikipedians and non-Wikipedians which, from Burke’s findings, this would indicate that Wikipedians are much more likely to succeed in a bid to become a Wikipedia administrator. This is a positive finding as administrators are involved more actively in the development and enforcement of policy and it seems logical that the people you would want to see in that role in an online community are people who are more active in the community as a whole.

The tie between communication and the invocation of norms and becoming a community enforcer may also hold true in other online communities. If so it would prove useful to be able to identify potential leaders/enforcers early on in order to make sure they remain in the community and rise to positions that suit them.

5. SUMMARY AND FUTURE WORK

Wikipedians are the essential core of the Wikipedia community. Previous work has revealed the volume of contributions to Wikipedia by Wikipedians and its prominence the work of Wikipedians has for readers. In this research we quantitatively analyze the qualitative assertions

of previous work and confirm that there is a relatively small group of editors on Wikipedia who produce more volume and higher quality edits. We also find that this core of editors is active about 30% of the time.

Wikipedians are born, not made. We have found that the initial activity of Wikipedians set them apart from the majority of other editors. This result suggests that casual users are rarely recruited to be more frequent editors. Future research could examine the effect that a welcoming process or mentoring program could have on recruiting more frequent contributors to cross the Wikipedia threshold.

Wikipedians are consistent. After performing at a high level from the beginning of their membership in Wikipedia, Wikipedians tend to maintain a high and constant level of participation for the majority of their lifespan. This result suggests that the natural intrinsic incentive system which is a part of that activities of editors is effective at maintaining the prolific editors.

Wikipedians don’t do more over time. With the exception of invoking community norms to explain their edits, Wikipedians do not do more work, better work, or more community-oriented work over time.

In addition to these findings, we have suggested several implications that follow from them: channeling new editors’ enthusiasm to accomplished necessary work, teach desired skills, and improve retention, crafting personalized appeals to solicit one-time editors to return, and promoting the use of Talk pages. We also identified important areas for future research, including alternate definitions of quality editing.

6. ACKNOWLEDGEMENTS

Thanks to the reviewers who took the time and effort to provide significant substantive feedback which has been incorporated. This work would not have been possible without the support and assistance of our research group. This material is based upon work supported under a National Science Foundation Graduate Research Fellowship. This work has also been supported by the National Science Foundation under grants IIS 05-34420 and IIS 03-24851.

7. REFERENCES

- [1] ADLER, B., AND DE ALFARO, L. A Content-Driven Reputation System for the Wikipedia. In *Proc. of WWW '07* (2007), vol. 7, pp. 261–270.
- [2] ALMEIDA, R., MOZAFARI, B., AND CHO, J. On the Evolution of Wikipedia. *Proc. of ICWSM 7* (2007).
- [3] BEENEN, G., LING, K., WANG, X., CHANG, K., FRANKOWSKI, D., RESNICK, P., AND KRAUT, R. Using Social Psychology to Motivate Contributions to Online Communities. In *Proc. of CSCW '04* (2004), ACM Press, NY, USA, pp. 212–221.
- [4] BRYANT, S. L., FORTE, A., AND BRUCKMAN, A. Becoming Wikipedian: Transformation of Participation in a Collaborative Online Encyclopedia. In *Proc. of GROUP '05* (2005), ACM Press, NY, USA, pp. 1–10.
- [5] BURKE, M., AND KRAUT, R. Taking up the Mop: Identifying Future Wikipedia Administrators. In *Proc. of CHI '08: Extended Abstracts* (2008), ACM Press, NY, USA, pp. 3441–3446.
- [6] BUTLER, B., JOYCE, E., AND PIKE, J. Don’t Look Now, But We’ve Created a Bureaucracy: The Nature

- and Roles of Policies and Rules in Wikipedia. In *Proc. of CHI '08* (2008), ACM Press, NY, USA, pp. 1101–1110.
- [7] COSLEY, D., FRANKOWSKI, D., KIESLER, S., TERVEEN, L., AND RIEDL, J. How Oversight Improves Member-Maintained Communities. In *Proc. of CHI '05* (2005), ACM Press, NY, USA, pp. 11–20.
- [8] COSLEY, D., FRANKOWSKI, D., TERVEEN, L., AND RIEDL, J. Using Intelligent Task Routing and Contribution Review to Help Communities Build Artifacts of Lasting Value. In *Proc. of CHI '06* (2006), ACM Press, NY, USA, pp. 1037–1046.
- [9] COSLEY, D., FRANKOWSKI, D., TERVEEN, L., AND RIEDL, J. Suggestbot: Using Intelligent Task Routing to Help People Find Work in Wikipedia. In *Proc. of IUI '07* (2007), ACM Press, NY, USA, pp. 32–41.
- [10] CUMMINGS, J., BUTLER, B., AND KRAUT, R. The Quality of Online Social Relationships. *Communications of the ACM* 45, 7 (2002), pp. 103–108.
- [11] DUCHENEAUT, N. Socialization in an Open Source Software Community: A Socio-Technical Analysis. In *Proc. of CSCW '05*, 4 (2005), ACM Press, NY, USA, pp. 323–368.
- [12] DUGAN, C., GEYER, W., MULLER, M., DiMICCO, J., BROWNHOLTZ, B., AND MILLEN, D. It's All 'About You': Diversity in Online Profiles. In *Proc. of CSCW '08* (2008), ACM Press, NY, USA, pp. 703–706.
- [13] EBNER, M., AND HOLZINGER, A. Lurking: An Underestimated Human-Computer Phenomenon. *IEEE MULTIMEDIA* (2005), pp. 70–75.
- [14] FISHER, D., SMITH, M., AND WELSER, H. You Are Who You Talk To: Detecting Roles in Usenet Newsgroups. In *Proc. of HICSS'06* (2006), vol. 3.
- [15] GILES, J. Internet Encyclopedias Go Head to Head. *Nature* 438 (2005), pp. 900–901.
- [16] HARPER, F., FRANKOWSKI, D., DRENNER, S., REN, Y., KIESLER, S., TERVEEN, L., KRAUT, R., AND RIEDL, J. Talk Amongst Yourselves: Inviting Users to Participate in Online Conversations. In *Proc. of the 12th international conference on Intelligent user interfaces* (2007), ACM Press, NY, USA, pp. 62–71.
- [17] HOGG, T., WILKINSON, D., SZABO, G., AND BRZOZOWSKI, M. Multiple Relationship Types in Online Communities and Social Networks. In *Proc. of the AAAI Symposium on Social Information Processing* (2008), pp. 30–35.
- [18] KARAU, S., AND WILLIAMS, K. Social Loafing: A Meta-Analytic Review and Theoretical Integration. *Journal Of Personality And Social Psychology* 65 (1993), pp. 681–681.
- [19] KITTUR, A., CHI, E., PENDLETON, B., SUH, B., AND MYTKOWICZ, T. Power of the Few vs. Wisdom of the Crowd: Wikipedia and the Rise of the Bourgeoisie. *alt.CHI '07* (2007).
- [20] KITTUR, A., AND KRAUT, R. Harnessing the Wisdom of Crowds in Wikipedia: Quality Through Coordination. In *Proc. of CSCW '08* (2008), ACM Press, NY, USA, pp. 37–46.
- [21] KITTUR, A., SUH, B., PENDLETON, B. A., AND CHI, E. H. He Says, She Says: Conflict and Coordination in Wikipedia. In *Proc. of CHI '07* (2007), ACM Press, NY, USA, pp. 453–462.
- [22] KRIPLEAN, T., BESCHASTNIKH, I., AND McDONALD, D. Articulations of Wikiwork: Uncovering Valued Work in Wikipedia Through Barnstars. In *Proc. of CSCW '08* (2008), ACM Press, NY, USA, pp. 47–56.
- [23] KUMAR, R., NOVAK, J., AND TOMKINS, A. Structure and Evolution of Online Social Networks. In *Proc. of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2006), ACM Press, NY, USA, pp. 611–617.
- [24] LAMPE, C., AND RESNICK, P. Slash (dot) and Burn: Distributed Moderation in a Large Online Conversation Space. In *Proc. of CHI '04* (2004), ACM Press, NY, USA, pp. 543–550.
- [25] LING, K., BEENEN, G., LUDFORD, P., WANG, X., CHANG, K., LI, X., COSLEY, D., FRANKOWSKI, D., TERVEEN, L., RASHID, A., ET AL. Using Social Psychology to Motivate Contributions to Online Communities. *Journal of Computer-Mediated Communication* 10, 4 (2005).
- [26] MARKS, K. Power Laws and Blogs. <http://homepage.mac.com/kevinmarks/powerlaws.html>, 2003.
- [27] ORTEGA, F., AND BARAHONA, J. M. G. Quantitative Analysis of the Wikipedia Community of Users. In *Proc. of WikiSym '07* (2007), ACM Press, NY, USA, pp. 75–86.
- [28] PRIEDHORSKY, R., CHEN, J., LAM, S. T. K., PANCIERA, K., TERVEEN, L., AND RIEDL, J. Creating, Destroying, and Restoring Value in Wikipedia. In *Proc. of GROUP '07* (2007), ACM Press, NY, USA, pp. 259–268.
- [29] SUH, B., CHI, E., PENDLETON, B., AND KITTUR, A. Us vs. Them: Understanding Social Dynamics in Wikipedia with Revert Graph Visualizations. In *Proc. of IEEE VAST '07* (2007), pp. 163–170.
- [30] SWARTZ, A. Who Writes Wikipedia? <http://www.aaronsw.com/weblog/whowritesWikipedia>, September 2006.
- [31] VIÉGAS, F., WATTENBERG, M., AND DAVE, K. Studying Cooperation and Conflict Between Authors with History Flow Visualizations. In *Proc. of CHI '04* (2004), ACM Press, NY, USA, pp. 575–582.
- [32] VIÉGAS, F., WATTENBERG, M., KRISS, J., AND VAN HAM, F. Talk Before You Type: Coordination in Wikipedia. In *Proc. of HICSS '07* (2007), IEEE.
- [33] VOSS, J. Measuring Wikipedia. In *Proc. of the 10th International Conference of the International Society for Scientometrics and Informetrics* (2005).
- [34] WHITTAKER, S., TERVEEN, L., HILL, W., AND CHERNY, L. The Dynamics of Mass Interaction. In *Proc. of CSCW '98* (1998), ACM Press, NY, USA, pp. 257–264.
- [35] WILKINSON, D., AND HUBERMAN, B. Assessing the Value of Cooperation in Wikipedia. *First Monday* 12, 4 (2007).
- [36] WILKINSON, D. M., AND HUBERMAN, B. A. Cooperation and Quality in Wikipedia. In *Proc. of WikiSym '07* (2007), ACM Press, NY, USA, pp. 157–164.