Cooperation and quality in Wikipedia

Dennis Wilkinson Bernardo Huberman
Information Dynamics Laboratory, Hewlett-Packard Labs
{dennis.wilkinson,bernardo.huberman}@hp.com

Abstract

The rise of the Internet has enabled collaboration and cooperation on an unprecedentedly large scale. The online encyclopedia Wikipedia, which presently comprises 7.2 million articles created by 7.04 million distinct editors, provides a consummate example. We examined all 50 million edits made to the 1.5 million English-language Wikipedia articles and found that the high-quality articles are distinguished by a marked increase in number of edits, number of editors, and intensity of cooperative behavior, as compared to other articles of similar visibility and age. This is significant because in other domains, fruitful cooperation has proven to be difficult to sustain as the size of the collaboration increases. Furthermore, in spite of the vagaries of human behavior, we show that Wikipedia articles accrete edits according to a simple stochastic mechanism in which edits beget edits. Topics of high interest or relevance are thus naturally brought to the forefront of quality.

Categories and Subject Descriptors H.5.3 [*HCI*]: Group and Organization Interfaces—Theory and models, Evaluation/methodology, Web-based interaction; K.4.3 [*Computers and Society*]: Organizational Impacts—Computer-supported collaborative work

General Terms Collaborative authoring, groupware

Keywords Cooperation, Wikipedia

WikiSym'07, October 21–23, 2007, Montréal, Québec, Canada.

Introduction

The online encyclopedia Wikipedia[1] provides an unprecedented example of large-scale, worldwide collaboration. Its 7.2 million articles have been generated from 282 million edits by 7.04 million distinct contributors[2], as of this writing. Wikipedia's exponential [33] growth since its inception in January 2001 has been enabled by the wiki interface [18], which allows any user to easily modify any article or to create new articles. This arrangement virtually eliminates the barrier to contribution, paving the way for intense activity at uncertain cost to article quality and value.

While Wikipedia's overall quality is difficult to measure in comprehensive way, its content has unquestionably been deemed useful and relevant by the user community at large. Its website is the 10th most visited on the Internet[3], serving an average of 18925 requests per second[4].

In light of its popular success, the question of which Wikipedia articles are high-quality, and how these articles are created, is of interest. A number of methods for automatic assessment of article quality have been proposed. In [19], the number of edits and unique editors to an article were suggested as metrics for quality, but no justification was provided. Other characteristics such as factual accuracy [13, 12, 10], credibility [6], revert times [31], and formality of language [11] have been used to assess small samples of Wikipedia's articles and in some cases compare them to articles of traditional encyclopedias. It is doubtful that encyclopedia quality can be assessed using a single metric (e.g. [8]), but complex combinations of metrics [30] depend on

[1] http://wikipedia.org
[2] http://meta.wikimedia.org/wiki/List_of_Wikipedias
[3] http://www.alexa.com/
[4] http://hemlock.knams.wikimedia.org/~leon/stats/reqstats/reqstats-monthly.png

rather arbitrary parameter choices. A crucial element lacking from previous tests of metrics of article quality is any consideration of article popularity or relevance, which can clearly be expected to affect metrics such as number of edits, number of links, article length, number of images, and many others.

Wikipedia's dynamics have also been studied in the context of evolution and network physics [34], with the addition of new articles described by a time-dependent acceleration mechanism [27] or a somewhat puzzling preferential attachment model [5]. Other work has examined the evolution of editors' roles and contributions in Wikipedia's development [16]. A power law relation was claimed for the distribution of edits per article [4, 33], but no mechanism was proposed and our data disagree with this assessment.

While this previous work contributes to the understanding of Wikipedia as a complex system, it does not provide insight into the development of quality at the level of individual articles, in particular the effects of large-scale collaboration and cooperation. While the potential benefits of cooperation are great [7], it is well-known from research in domains such as software development (e.g., [3]) and industrial design (e.g., [1]) that increasing the size of a collaboration can have an ambiguous or even deleterious effect on result quality [14]. In Wikipedia, moreover, there is no selection process for editors (beyond having a computer, an internet connection, and the time and energy to devote to unpaid work) and no direct supervision of editing; and the "distributed moderation" system used in Wikipedia to resolve disputes has been shown in other domains to be slow and only partially effective [17]. While the wiki interface does facilitate coordination [9, 20, 21], it cannot resolve cultural [22] or philosophical differences which result in a significant number of "edit wars" and mass deletions [31].

One might therefore suspect the best articles in Wikipedia to be produced by a few dedicated, exceptional editors. Indeed, this view has been recently espoused by some of the most influential members of the Wikipedia community [25].

In this paper we demonstrate that there is a strong overall correlation between number of edits, number of distinct editors, and article quality. This correlation is observed after article visibility, popularity, and age have been carefully taken into account. We further exhibit evidence of more cooperation in the development of the high-quality articles than other articles. This evidence includes a strong correlation between discussion (talkpage) activity and article quality, more edits per editor to high-quality articles, and a markedly different pattern of editors' responses to other edits on these pages. Our study encompasses all 50 million non-robot edits to the 1.5 million articles in the English language Wikipedia made between its inception in January 2001 and November 2, 2006.

We also show that Wikipedia articles accrete edits according to a simple feedback mechanism in which edits beget edits. That is, the number of new edits to a given article in a given period of time is a randomly varying percentage of the total number of previous edits. This stochastic process produces a lognormal distribution in the number of edits per article for articles created during a particular time slice, where the distribution parameters μ and σ^2 depend linearly on the age of the time slice. A statistical test is shown to be in strong agreement with the lognormal distribution, and the linear evolution in time of μ and σ^2 is demonstrated.

The lognormal distribution of edits per article in Wikipedia means that a small but significant population of articles experience a disproportionately high number of edits and editors, while the vast majority of articles undergo far less activity. The heavy tail of highly-edited articles, representing topics of particular interest or relevance, are thus naturally brought to the forefront of quality. In addition, the mechanism implies that articles do not reach a steady state but continue to accrete edits regardless of their age, a prediction which is confirmed by the data.

The structure of this paper is as follows. We first discuss the overall dynamics of the growth of Wikipedia articles and their implications. We then demonstrate the correlation between quality and increased editing activity. Finally, we present evidence for increased cooperation in the high-quality articles. Our data set, including the method used to remove edits made by robots, are described in the appendix. The statistical tests used to justify the lognormal distribution of edits per article are also described in the appendix.

Dynamics of article growth

To address the subject of how edits contribute to article quality, we first examine underlying patterns in the way Wikipedia articles accrete edits. While individual users exhibit highly variable editing activity, the overall dynamics of edit accretion is well-described by the following simple mechanism.

Consider the number of new edits $\Delta n(t)$ to an article made between time t and time $t + dt$, an interval of perhaps several hours. Of course, complicated fluctuations in human behavior and activity cause this number to vary in a random way, but we claim that $\Delta n(t)$ is on average proportional to the total number of previous edits. This is expressed mathematically as

$$\Delta n(t) = [a + \xi(t)]n(t),$$

where $n(t)$ is the total number of edits to a given article up until time t, a is a constant (average) rate of edit accretion, and $\xi(t)$ is mean-zero random term accounting for fluctuations. The total number of edits at time $t+dt$ is thus given by

$$n(t+dt) = n(t)+\Delta n(t) = [1+(a+\xi(t))]n(t). \quad (1)$$

Because of the random nature of human activity embodied by $\xi(t)$, the number of edits to a given article at a given time can be predicted only within a range of values specified by a probability distribution. Previous work on similar processes, such as the evolution of the World Wide Web [15] and many others (e.g., [26]), has shown that the distribution resulting from equation (1) is lognormal[5] and given by

$$P[n(t)] = \frac{1}{n\sqrt{2\pi}\sqrt{s^2 t}} \exp\left[-\frac{(\log n - at)^2}{2(s^2 t)}\right], \quad (2)$$

where s^2 is the variance of the $\xi(t)$. This equation shows that the distribution parameters $\mu = at$ and $\sigma^2 = s^2 t$ are linearly related to the age t of the article. μ and σ^2 represent the mean and variance, respectively, of the log of the data, and are thus related to but not equal to the distribution mean and variance. In practice, we considered articles created during a time slice of average age t in order to obtain enough data points to constitute a distribution. Provided the time slice is not too long, editing within the slice does not corrupt the distribution much.

Equation (2) was verified by a study of the 50.0 million edits made by the 4.79 million non-robot contributors to the 1.48 million articles of the English-language Wikipedia between its inception in January 2001 and November 2, 2006. A statistical test of all time slices yields a p-value of greater than 0.5 for 47.8 % of the 3688 relevant time slices for the lognormal distribution (further details on the test and the data are provided in the appendix).

The distribution of edits for articles in various time slices is illustrated in figures 1, showing the actual lognormal distribution with its heavy tail, and 2, which more effectively demonstrates the accuracy of the prediction of the model. Note in figure 2 that the distribution mean increases with age, as expected from the feedback model. The overall growth of Wikipedia is also observed in the form of higher total numbers of younger articles.

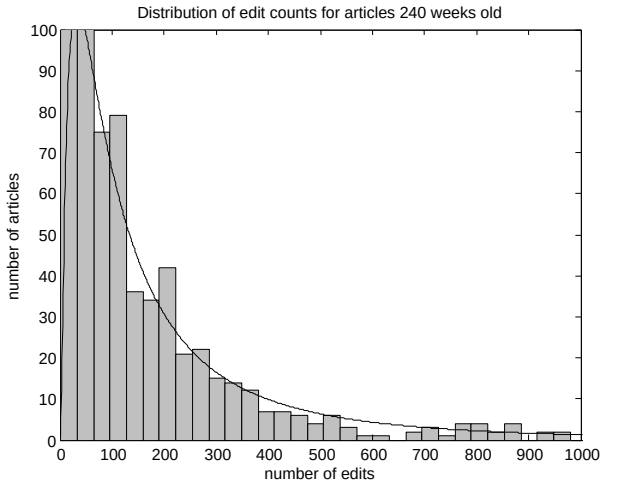

Figure 1. Distribution of the number of edits per article for articles of age $t = 240$ weeks. The plot was truncated at the high end of both axes for readability; in fact, there are articles in this set with many thousands of edits, and hundreds of articles with very few edits. The best fit lognormal curve is included for comparison.

[5] In equation 1, the noise terms at different t are assumed to be uncorrelated. In fact, as one might expect, the percentage increase in edits does demonstrate a small positive autocorrelation over periods of less than 20 to 30 days. Since the autocorrelation length is finite, however, the central limit theorem may still be applied to obtain a log-normal distribution; the difference is that the rate parameter a must be modified to account for the autocorrelation [2]. Because the modification is small, for the sake of simplicity, we do not include it here.

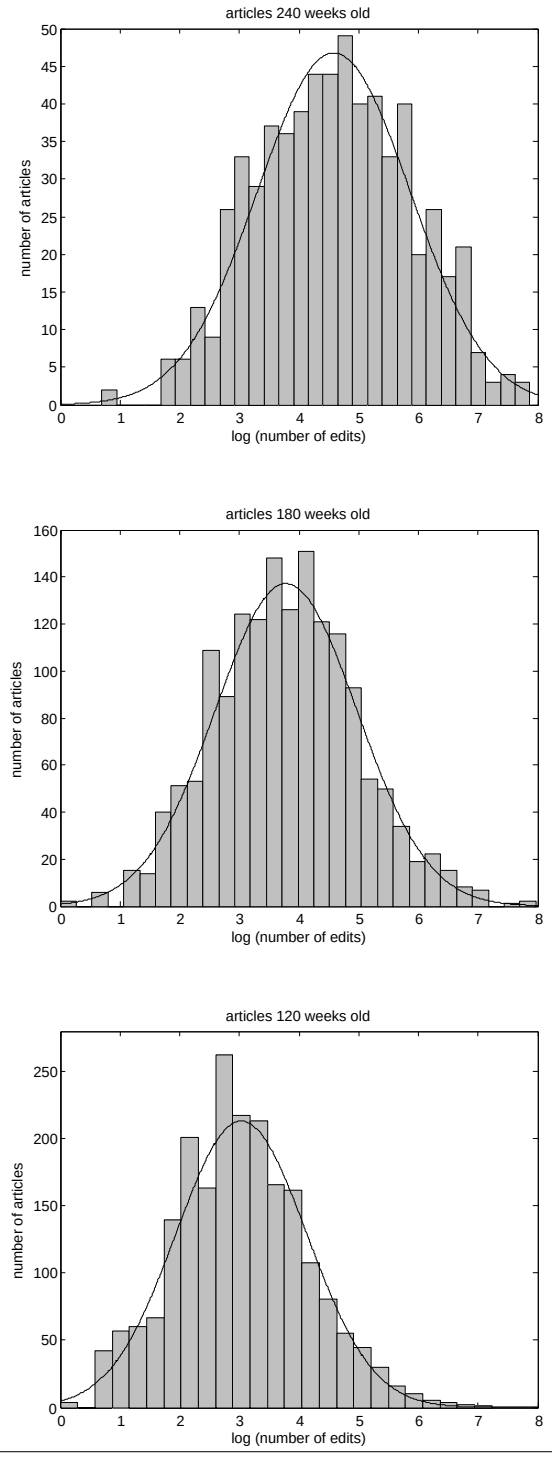

Figure 2. Distributions of the logarithm of the number of edits per article for articles of ages t = 240 (top panel), 180 (middle), and 120 (bottom) weeks. Because the distribution of edits per article is lognormal, the logarithm is normally distributed, and the best fit normal curve is included for comparison.

The variation of the distribution parameters μ and σ^2 with age is demonstrated in figure 3. The linear dependence is highlighted by the fitted curve[6]. Anomalous time slices which do not fit the overall trend include two periods in which a large number of rather trivial articles with low edit counts were created at once, and the recent data containing a large number of short "stub" articles which have yet to be combined into regular articles or deleted. These slices contain an unusually high number of articles with low edit counts.

The lognormal distribution has a heavy tail at the high end, implying that a small number of articles accrete a disproportionally large number of edits. As we show below, edits correspond on average to an increase in article quality. The feedback mechanism of edit accretion thus results in a small body of high quality articles. These high quality articles deal with topics of high visibility or relevance, while the vast majority of Wikipedia articles are relatively infrequently edited and have far lower visibility[7].

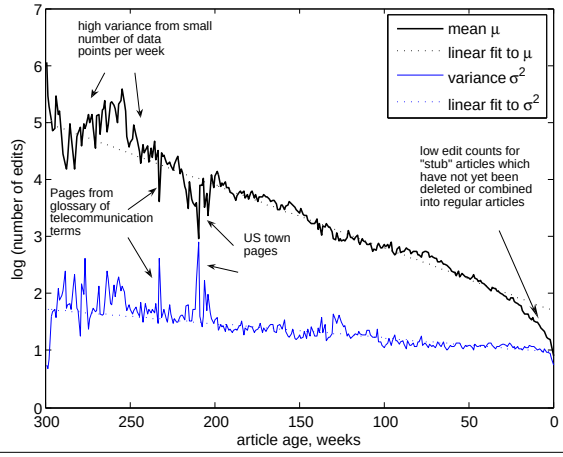

Figure 3. Evolution of the mean μ (top) and variance σ^2 (bottom) of the lognormal distribution of edits per article. The articles are grouped by age in weeks.

[6] The reason that the fitted line does not go to 0 for as the age approaches zero is the correction due to the finite autocorrelation time in the $\xi(t)$ term discussed in the previous footnote.

[7] Since each time slice follows a lognormal distribution, the overall distribution of edits per article is a mixture over time of lognormals with linearly varying μ and σ^2, multiplied by a factor accounting for the overall growth of Wikipedia. This integral is not analytic for the time frame of our data, but by numerical integration can be shown to retain a lognormal character. In the long-time limit, the integral tends towards a power law if overall growth is exponential [15]. However, the time frame of our data set is not the long time limit, in contrast to the findings of [4] who report a power law fit for the overall distribution of edits per article.

Editing and article quality

As discussed in the introduction, it is of interest to determine whether extensive editing by a large number of diverse contributors increases article quality.

To test for a correlation between editing and article quality, we compared counts of edits and distinct editors on Wikipedia "featured" articles to corresponding counts all other articles. Featured articles are selected by the Wikipedia community as "the best articles in Wikipedia," according to criteria such as accuracy, neutrality, completeness, and style[8]. They undergo a rigorous review process and are demoted if they do not continue to uphold these high standards. In our data set, there were 1211 featured articles, or 0.081 % of all English-language articles.

It is crucially important to control for topic popularity, article visibility, and article age when comparing editing activity on two populations of articles. If one of the populations has, for example, more popular articles than the other, then it may have more edits per article simple because of popularity. The connection between age and editing was demonstrated in figure 3, while the connection between popularity or visibility and editing is apparent, but not the main point of, figure 4 below. To control for these factors, we took the following three steps.

First, to account for topic popularity, we noted that Google pagerank[9] has been shown to correlate strongly to the number of times a Wikipedia page is viewed [28] and is thus a useful and accurate proxy. We grouped the articles by their pagerank, an integer between 0 and 9, and compared the featured and non-featured articles separately for each pagerank.

Second, we normalized the counts of number of edits and distinct editors by article age before comparing articles of different populations. In particular, for an article A of age t having undergone n edits, we computed the quantity

$$x(A) = \frac{\log n - \mu(t)}{\sigma(t)}, \qquad (3)$$

where $\mu(t)$ and $\sigma(t)$ were previously measured to be the average and standard deviation, respectively, of $\log n$ for all articles of age t. The logarithm of n was used because, as previously discussed, the overall distribution of edit counts on Wikipedia articles is roughly lognormal[10]. An analogous procedure was used to compute an age-normalized measure of the number of distinct editors for each article.

Finally, to account for the special attention featured articles may receive while they are mentioned on the main page and while they are under review for featured status, we removed edits generated during the two most active weeks for each article. In fact, this is probably not necessary, since the percentage of edits made during the two most active weeks for the featured population (13.2 %) is actually far lower than for the rest of the articles (19.2 %), or if age is taken into account, almost equivalent. The same is true if periods of one week, three weeks and one month are considered.

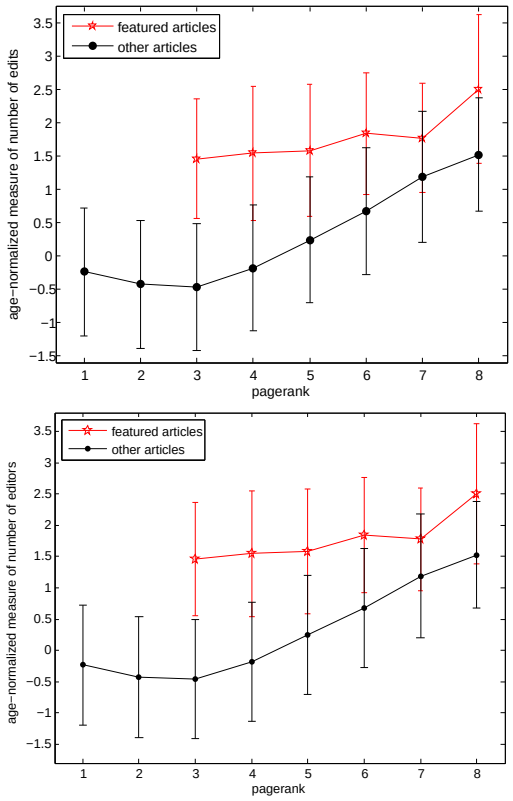

Figure 4. Average and standard deviation (error bars) of the age-normalized measure of number of edits (top, equation 3), and number of distinct editors (bottom, analogous expression), grouped by pagerank. For each pagerank, the discrepancy between featured and non-featured populations is statistically significant by a huge margin.

[8] http://en.wikipedia.org/wiki/Wikipedia:Featured_articles

[9] http://www.google.com/technology/

[10] The mean and variance of samples from a lognormal distribution are overly sensitive to variations because of the heavy tail, whereas $\log n$ is roughly normally distributed and its samples will be very faithful to the true distribution.

The comparison of the number of edits and distinct editors for featured and nonfeatured populations is shown in figure 4, demonstrating the correlation between number of edits, number of distinct editors, and article quality in Wikipedia. Because of the large number of articles in the data set, the difference between the populations for each pagerank are statistically significant by an exceedingly wide margin[11].

As to the question of causality between edits and quality, recall that in general, as we showed, articles continue to accrete edits and evolve instead of reaching a steady state. Resolving causality in an aggregate sense is thus most likely impossible. Indeed, the development of an article is a highly complex process [29] and both directions of causality between editing and quality are likely to play a role.

Cooperation and article quality

As a final subject, we consider the question of how Wikipedia editors collaborate to produce high-quality articles, as opposed to other articles. While detailed examination has been used to explore the evolution of individual articles [31], such methods are difficult to apply on a large scale. In this section we again consider all edits to 1.5 million articles in the English-language Wikipedia and compare metrics of cooperation on the featured and nonfeatured populations.

Most Wikipedia articles have an associated talkpage[12], which editors use to engage in discussions about the content of the article. One measure of cooperation for an article is the amount of activity on its corresponding talkpage. To measure this activity, we obtained the number of revisions to talkpages and compared the populations of featured and nonfeatured populations, following the same procedure as for articles. The results of this procedure are shown in figure 5 and demonstrate a strong correlation between number of comments posted to a talkpage and quality of the corresponding article. As for the articles, the differences between the populations are statistically significant by a wide margin. It is worth noting that the difference between the featured and nonfeatured populations is more distinct in this figure than the corresponding plots for

[11] For example, the probability (p-value) that the featured and nonfeatured pagerank 7 articles come from the same distribution is uncalculably small: less than 10^{-16}. A t-test is appropriate because the statistic x from equation 3 is approximately normally distributed. See also table in the appendix.

[12] http://en.wikipedia.org/wiki/Wikipedia:Talk_page

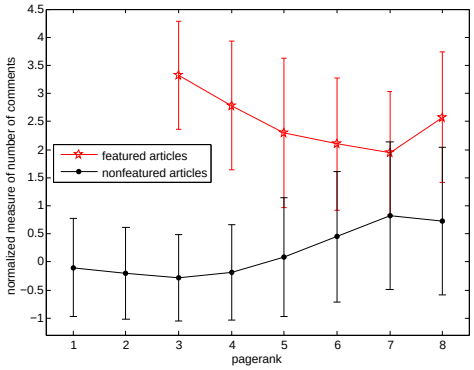

Figure 5. Average and standard deviation (error bars) of the age-normalized measure of number of revisions posted to talkpages, grouped by the corresponding article's pagerank.

edits and distinct editors, suggesting that cooperation could be a more important indicator of article quality than raw edit counts. It is also interesting that talkpage revisions decrease with pagerank within the featured population.

Other macroscopic measurements of the featured and nonfeatured populations provide further evidence that cooperation is correlated to article quality. Figure 6 shows the number of edits per editor for the two populations. An editor is very unlikely to engage in cooperative authoring without making at least several edits. The greater number of edits per editor for the high-quality population indicates that such behavior is more common in this set. By contrast, popularity of topic is almost completely uncorrelated to editor dedication, on average, as shown by the nearly constant number of edits per editor over all pageranks for the nonfeatured population.

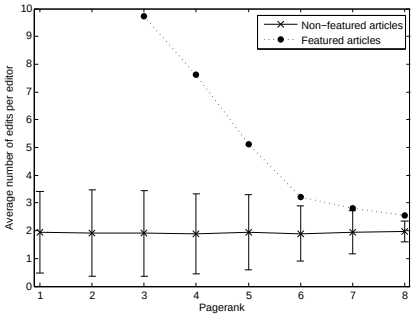

Figure 6. Number of edits per editor for nonfeatured (average and standard deviation) and featured (average) populations, grouped by pagerank.

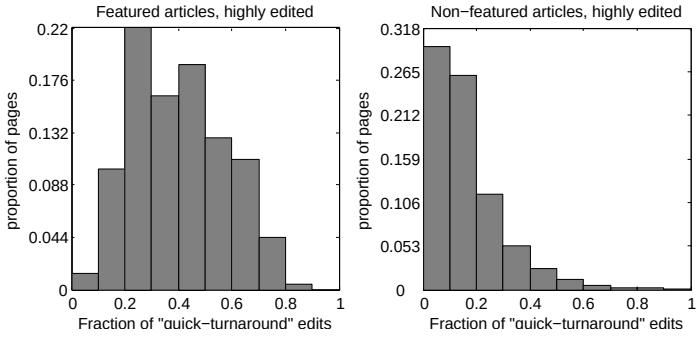

Figure 7. Proportion of highly-edited articles, grouped by percentage of times an edit followed within 30 minutes of the previous edit and was made by a different (human) editor

Finally, figure 7 shows proportions of highly-edited articles grouped by the percentage of time an edit was made within 30 minutes of the previous edit and by a different (human) editor. Articles were chosen as highly-edited if their age-normalized measure of the logarithm of the edit counts (equation 3) was greater than 2.5; there were 229 such featured articles and 19056 nonfeatured. The period of 30 minutes was chosed arbitrarily, but a similar trend is observed for periods ranging from 10 minutes to an hour. Similar results are also observed for less highly-edited populations of articles.

The different trends exhibited in the panels of figure 7 for the featured and non-featured populations suggest that, in general, the article creation process occurs differently for these two populations. Edits which are not "quick succession" are made either long after the previous edit, or by the same editor who made the previous edit. When most of an article's edits were not made in quick succession, it indicates that the article was created primarily by individuals working alone. The larger percentages of quick succession edits in the high-quality population of articles thus suggests that the creation process involved more cooperation[13].

[13] Increased cycles of vandalism and repair or "edit wars" [31] could also play a role in the trends of figure 7, but such analysis is beyond the scope of this paper

Conclusion

We have shown that the high-quality articles in Wikipedia are distinguished from the rest by a larger number of edits and distinct editors, having carefully controlled for article visibility, popularity, and age. Furthermore, we demonstrated more intense patterns of cooperation in the high-quality articles than in other articles. These findings are in contrast to observations of cooperative efforts in other domains where result quality does not necessarily increase with the number of collaborators. While we did not explore the question of how Wikipedia succeeds where other large collaborative ventures fail, possible reasons include the efficiency of the wiki interface, the Wikipedia community's strong emphasis on coordination and organization [32], and details of the processes and policies used to facilitate cooperation [23].

Additionally, we have have shown that although Wikipedia is a complex system in which of millions of individually unpredictable editors collaborate in an unscheduled and virtually uncontrolled fashion, article growth follows a very simple overall pattern on average. This pattern implies that a small number of articles, corresponding to topics of high relevance or visibility, accrete a disproportionately large number of edits, while the vast majority of articles experience far less activity. Subjects of particular importance or popularity are thus naturally brought to the forefront of quality, validating Wikipedia as a successful collaborative effort.

Acknowledgments: We thank Travis Kriplean for his work in helping process the data set and Yuri Karaban for his Perl module. A preliminary version of this paper appeared in *First Monday* **12**, 4 (2007).

Appendix: Methods and Data

The raw data for our study were all 55.3 million edits to the English-language Wikipedia made between Wikipedia's inception in January 2001 and November 2, 2006. This data included username or url, page title, and timestamp. From the raw data, we eliminated redirect and disambiguation pages, which are articles with no content that merely point to other articles, and edits made by robots. Redirects and disambiguation pages were identified using simple text analysis. Robot edits were identified using the list of registered Wikipedia robots[14], and by identifying edits made by a single user in improbably quick succession. This process eliminated 5.23 million edits, or 9.5 % of the original 55.3 million.

A small percentage of articles were not used because of technical difficulties in the title caused by rare foreign characters. Google

[14] http://en.wikipedia.org/wiki/Wikipedia:Registered_bots

pageranks were obtained by querying Google[15]. Some recent articles did not yet have a pagerank and so pagerank zero articles were not included in the analysis.

To test the lognormal fit and obtain the quoted p-value, we applied a typical χ^2 fitting procedure to each time slice using the likelihood ratio statistic [24]. In this test, the time slice length was variable because of the overall growth of Wikipedia; more recent articles were grouped into smaller slices because otherwise the distribution was skewed by edits made within the slice. In practice, we chose time slices long enough to contain 400 articles. The expected distribution for each slice was calculated using the slice's sample mean and variance, and the data was grouped into bins whose width was the minimum required to make the expected count greater than 8. Of course, slight variations in the quoted p-value, on the order of several percent, were obtained by varying the time slice length and bin size.

Finally, the following tables provide justification that article age and popularity are different for the featured and nonfeatured populations.

pagerank	≤ 3	4	5	6	7	≥ 8
featured	0.4	3.3	18.9	29.8	46.3	1.4
nonfeatured	10.7	47.7	34.4	6.7	0.4	0.01

Table 1. Percentages of articles, by pagerank, for featured and nonfeatured populations

pagerank	≤ 3	4	5	6	7	≥ 8
featured	440	564	823	1310	1360	1650
nonfeatured	364	486	661	1150	1570	1701

Table 2. Average age, in days, of article populations

References

[1] T. J. Allen. Studies of the problem solving process in engineering. *IEEE Trans. Engineering Management*, EM-13(2):72–83, 1966.

[2] K. N. Berk. A central limit theorem for m-dependent random variables with unbounded m. *Ann. Prob.*, 1(2):352–354, 1973.

[3] F. Brooks. *The Mythical Man-month*. Addison-Wesley, Reading, Mass., 1975.

[4] L. Buriol, C. Castillo, D. Donato, S. Leonardi, and S. Millozzi. Temporal evolution of the Wikigraph. In *Proc. of the Web Intelligence Conference*, Hong Kong, 2006.

[5] A. Capocci, V. Servidio, F. Colaiori, L. Buriol, D. Donato, S. Leonardi, and G. Caldarelli. Preferential attachment in the growth of social networks: the case of Wikipedia. *Phys. Rev. E*, 74:036116, 2006.

[6] T. Chesney. An empirical examination of Wikipedias credibility. *First Monday*, 11(11), 2006.

[7] Scott Clearwater, Bernardo Huberman, and Tad Hogg. Cooperative solution of constraint satisfaction problems. *Science*, 254(5035):1181–1183, 1991.

[8] H. Crawford. Encyclopedias. In R. Bopp and L. C. Smith, editors, *Reference and information services: an introduction, 3rd ed.*, pages 433–459, Englewood, CO, 2001. Libraries Unlimited.

[15] Yuri Karaban's PERL module was very helpful for this: http://search.cpan.org/∼ykar/WWW-Google-PageRank-0.12/

[9] A. Desilets, S. Paquet, and N. Vinson. Are wikis usable? In *Proc. ACM Wikisym*, 2006.

[10] Editorial. Britannica attacks... and we respond. *Nature*, 440, 2006.

[11] W. Emigh and S. Herring. Collaborative authoring on the Web. In *Proc. HICSS*, 2005.

[12] *Encyclopaedia Britannica*. Fatally flawed: refuting the recent study on encyclopedic accuracy by the journal Nature, March 2006.

[13] Jim Giles. Internet encyclopaedias go head to head. *Nature*, 438:900–901, 2005.

[14] N. Glance and B. A. Huberman. The dynamics of social dilemmas. *Scientific American*, March 1994.

[15] B. A. Huberman and L. Adamic. Growth dynamics of the World Wide Web. *Nature*, 401:131, 1999.

[16] A. Kittur, E. Chi, A. Pendleton, B. Suh, and T. Mytkowicz. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. submitted for publication.

[17] C. Lampe and P. Resnick. Slash(dot) and burn: distributed moderation in a large online conversation space. In *Proc. ACM CHI Conf. Human Factors in Comp. Systems*, 2004.

[18] B. Leuf and W. Cunningham. *The Wiki Way*. Addison-Wesley, Boston, 2001.

[19] A. Lih. Wikipedia as participatory journalism. In *Proc. 5th International Symposium on Online Journalism*, austin, TX, 2004.

[20] E. Da Lio, L. Fraboni, and T. Leo. Twiki-based facilitation in a newly formed academic community of practice. In *Proc. ACM Wikisym*, 2006.

[21] A. Majchrzak, C. Wagner, and D. Yates. Corporate wiki users: results of a survey. In *Proc. ACM Wikisym*, 2006.

[22] U. Pfeil, P. Zaphiris, and C. S. Ang. Cultural differences in collaborative authoring of Wikipedia. *J. Computer-mediated Comm.*, 12(1), 2006.

[23] J. M. Reagle. A case of mutual aid: Wikipedia, politeness and perspective taking. In *Proc. Wikimania*, 2005.

[24] J. Rice. *Mathematical Statistics and Data Analysis*. Duxbury Press, Belmont, CA, 2nd edition, 1995.

[25] D. Riehle. How and why Wikipedia works: an interview. In *Proc. ACM Wikisym*, 2005.

[26] S. M. Ross. *Stochastic processes*. John Wiley, New York, 1996.

[27] D. Smith, J. Onnela, and N. Johnson. Accelerating networks. preprint.

[28] A. Spoerri. What is popular on Wikipedia and why? *First Monday*, 12(4), 2007.

[29] B. Stvilia, M. B. Twidale, L. C. Smith, and L. Gasser. Information quality work organization in wikipedia. preprint.

[30] B. Stvilia, M. B. Twidale, L. C. Smith, and L. Gasser. Assessing information quality of a community-based encyclopedia. In *Proc. ICIQ 2005*, pages 442–454, Cambridge, MA, 2005.

[31] F. B. Viégas, M. Wattenberg, and K. Dave. Studying cooperation and conflict between authors with history flow visualizations. In *Proc. SIGCHI*, 2004.

[32] F. B. Viégas, M. Wattenberg, J. Kriss, and F. van Ham. Talk before you type: coordination in Wikipedia. In *Proc. HICSS*, 2007.

[33] J. Voss. Measuring Wikipedia. In *Proc. ISSI 2005*, Stockholm, 2005.

[34] V. Zlotic, M. Bozicevic, H. Stefanic, and M. Domazet. Wikipedias: Collaborative web-based encylopedias as complex networks. *Phys. Rev. E*, 74, 2006.