

Mining for Practices in Community Collections: Finds From Simple Wikipedia

Matthijs den Besten¹, Alessandro Rossi², Loris Gaio², Max Loubser³ and Jean-Michel Dalle⁴

1 University of Oxford, Oxford e-Research Centre, 7 Keble Road, OX13QG, Oxford, UK
matthijs.denbesten@oerc.ox.ac.uk,
WWW home page: <http://users.ox.ac.uk/~ierc0002>

2 University of Trento, Department of Computer and Management Science, Via Inama 5, I-38100, Trento, Italy
{arossi,lgaio}@cs.unitn.it,
WWW home page: <http://rock.cs.unitn.it/{arossi,lgaio}>

3 Oxford Internet Institute, 1 St Giles, OX13JS, Oxford, UK
max.loubser@magd.ox.ac.uk,
WWW home page: <http://people.oii.ox.ac.uk/loubser>

4 Université Pierre et Marie Curie, c/o Agoranov, 3, rue Castex 75004, Paris, France
jean-michel.dalle@upmc.fr,
WWW home page: <http://www.upmc.fr>

Abstract. The challenges of commons based peer production are usually associated with the development of complex software projects such as Linux and Apache. But the case of open content production should not be treated as a trivial one. For instance, while the task of maintaining a collection of encyclopedic articles might seem negligible compared to the one of keeping together a software system with its many modules and interdependencies, it still poses quite demanding problems. In this paper, we describe the methods and practices adopted by Simple Wikipedia to keep its articles easy to read. Based on measurements of article readability and similarity, we conclude that while the mechanisms adopted by the community had some effect, in the long run more efforts and new practices might be necessary in order to maintain an acceptable level of readability in the Simple Wikipedia collection.

1 Introduction

Issues pertaining to how communities form, operate, and sustain themselves loom large in any project where voluntary participants gather to create something of value. There are growing concerns regarding which methods and practices within these communities account for their success. This is particularly important if one is interested in replicating successful models to other types of commons-based peer productions. This question is also fueled by quality concerns: if we are to employ a

peer produced collection, we would want to be reassured that what is made available is of sufficient quality and will remain so in the near future. The above considerations hold as much for open source communities developing systems like Linux and Apache as for open content communities creating collaborative repositories like Wikipedia.

Wikipedia is a large-scale collaborative project where voluntary contributors join forces to put together an encyclopaedia that is free for everyone to read and edit. The collaboration has produced, amongst others, an English version with over 2 million articles, a German version with over 600,000 articles and a Simple English version with over 20,000 articles. At the same time, Wikipedia has become extremely visible: for instance, its pages frequently appear high in the result list of search engines [2] and are nowadays commonly referenced by the media. In this paper we take a closer look at the organization of Wikipedia: we speculate on the problems it faces and on the way they are addressed, and, ultimately, we assess its potential for the future: will Wikipedia become an even more established source of information than it already is or will its efforts whither under the weight of vandalism and bureaucracy [7, 14]?

At first sight, Wikipedia is a clear success. On 15 December 2005, *Nature* published a special report comparing the quality of online encyclopaedias. Its findings highlighted that the upstart Wikipedia had a number of errors comparable to Encyclopaedia Britannica [9]. Britannica was compared to Wikipedia for quality by sampling 42 science articles that were sent to experts in the relative field for blind review. While the study was criticized by the Britannica organization for the way articles were selected and compared and clearly suffers from the weakness of a small sample from a large collection, it was influential in establishing that the distributed collaboration on Wikipedia can show performance approximately comparable to traditional methods of encyclopaedia production.¹

In terms of gross production, Wikipedia seems healthy too: content grew exponentially from 2002 to 2005 [18]. However, this is only a rough approximation of performance and does not take into account the quality of the content and as Denning *et al.* [4] note, a proper assessment of the performance and quality of Wiki projects would require attention to many more factors.

With tremendous growth, performance is affected by how well maintenance and organizational work is carried out and how well this non-article related work can be scaled. Reagle [15] examines the way the Wikipedia community develops a neutral point of view in articles and the norms that govern meta-discussions, such as those that take place on article discussion pages. Kittur *et al.* [10] measure directly the impact of these meta-discussions on the performance of the project. They found that the proportion of work on non-article content (such as discussion and fixing

¹ It is worth to note also that while all Wikipedia's inaccuracies raised by the Nature's study has been fixed in a matter of few days (the last being corrected on 25 January, 2006). On the contrary, at least 8 errors (out of the 64 inaccuracies that were not contested by Britannica) still stand nowadays untouched into the current online version of the Encyclopaedia Britannica.

vandalism) is increasing in Wikipedia, showing that more overhead is required as the collection grows and casting doubt over its sustainability.

The increasing demand for maintenance work includes the need to undo the actions of those with malicious intent. In this regard Wikipedia appears to be performing satisfactorily. Even in the early stages of development, it was clear that malicious edits were removed quickly: Ciffolilli [1] argues that it is the nature of the Wiki that makes it easy for friendly contributors to “clean up” vandalism. Later work by Viegas *et al.* [17] showed that “obscene” edits from their sample had a median lifetime of only 1.7 minutes.

Doubts remain over the long-term viability of Wikipedia as a reliable source of information [6, 11]. From the outside, vandalism and well-meaning ignorance are a continuous threat to article quality and from the inside, the dominance of difficult people with lots of time on their hands could easily become a danger of equal if not larger size. In the end, all depends on how well Wikipedia as a community manages to address these threats.

In the paper we take some steps towards studying how organizational practices enable to increase the overall quality of the collection of articles belonging to Simple Wikipedia, a spinoff project of the larger Wikipedia. In order to do so, we have to take into account that communities like Simple Wikipedia might employ methods and practices very different from the ones used in other commons based peer communities. In this sense, our work is very sympathetic with the theory of information goods production developed by Mateos Garcia and Steinmueller [13]. These scholars identify four distinct roles in the production process – that of the author, the editor, the integrator and the conductor – and argue that what distinguishes the “vertical” production of systems like Linux from the “horizontal” production of collections like Wikipedia, is that the cumulative dependencies of components in vertical production make the task of editor and integrator far more valuable and rewarding than the complementary dependencies in horizontal production. What we set out to do for the study reported in this paper was to see if we could identify integration efforts in collections anyway despite the low payoff for volunteers.

In what follows, we report on our findings. We first describe how we selected Simple Wikipedia as a collection, and why it is a good case study to address issues regarding quality and level of integration of the collection. We then lay out what data we have retrieved and how we have mined those data. Finally, we present our findings, which suggest that in the case of Simple Wikipedia that the integration efforts that we were able to identify have some effect but seem insufficient in order to maintain the quality of the collection in the long run.

2 Simple Wikipedia

In order to test our approach, we did a pilot study on Simple Wikipedia. Simple Wikipedia is a spin-off of Wikipedia that was initiated in 2003 because people felt

that many articles in Wikipedia were too hard to read, due to jargon, formality, or for other reasons – especially for children and non-native speakers of English. Simple Wikipedia, was the hope, would be the place where people go to look for an easily readable descriptions of topics. At the same time, contributors to Simple Wikipedia would commit to the ideals of this sub-project or at least adhere to an editorial policy that calls for greater readability when contributing descriptions of topics.

Simple Wikipedia is an ideal candidate to study the performance of commons based peer communities, since it is relatively easy to determine how well it adheres to its goals, zoom in on efforts that are made to address specific problems, and assess the result of these efforts. First of all, with less than 20,000 articles in its collection Simple Wikipedia is a relatively small encyclopedia project. Consequently, no extraordinary computational resources are needed to extract the project archive and analyze its contents. Besides, Simple Wikipedia is a project centered on a very specific goal: readability. What’s more, readability is something that can be measured. So, here we have a project that we can assess on its own terms. Moreover, Simple Wikipedia, as a separate project, was able to experiment and implement with its own editorial policies and managerial policies specifically to help attain its goal of simplicity. Most importantly, Simple Wikipedia has come to rely on the tag “unsimple” – more recently called “complex” – to single out articles, which do not meet its standards and need to be improved.

3 Data

The archive of Simple Wikipedia is available from the web at <http://downloads.wikimedia.org>. Our analysis here is based on the archive of July 2007, which contains the revision history of over 25,000 articles and around 27,000 pages of a different type such as discussion pages and user-pages where regular contributors present themselves. For each edit on an article, the archive lists the IP address or user-name and user-id of the editor, the time of edit, comments made by the editor, and the actual text of resulting from the edit. The text is marked up with tags to identify structural elements like sections and sub-sections and tags of a different type to identify labels – also called templates – that are applied to the text. For instance, an article that is considered to be hard to read will contain the string “`{{unsimple}}`” in the raw text. The status of an editor is slightly harder to determine. We can easily distinguish between editors who are known to the system and editors who contribute anonymously as the latter are identified by their user-id while for the former only the IP address is listed. Bots, scripts that carry out small repetitive edits such as spell checks and interlinking of articles, usually have a user-name that ends with “bot”. In addition, the user-id of these bots is listed as a user belonging to a special group in the auxiliary user-group table, as is the user-id of the users with special rights known as administrators.

3.1 Metrics

In the results that we outline further in this section, we make use of two classes of computational linguistics metrics as a way to characterize changes, made overtime in the Simple Wikipedia articles, resulting from the collaborative and iterative editing process. We are in particular interested in studying two different, albeit interconnected, evolving features of articles: first of all we want to assess how the readability of an article changes over time; subsequently we want to focus on change *tout court* by introducing measures of similarity, accounting for the degree of change/persistence of information in the various revisions of the articles.

The readability of an article is determined by computing the Flesch readability score of the article's text with help of the GNU Style package. This score is a function of the number of syllables per word and the number of words per sentence [8]. More precisely, the formula 'score = 206.835 - 84.6*syllables/words - 1.1015*words/sentences' yields a number that is usually between 0 and 100 and between 60 and 70 for standard English texts. This Flesch reading easy formula, which has been elaborated on the basis of school texts by Flesch in 1948, has been very popular, especially in the US, as a measure of plain English. Its popularity rests on the fact that the formula is easy to compute, yet often accurate. Even word processing applications, such as Microsoft Word, often provide the score as part of their statistics.

The second class of language metrics focuses more explicitly on characterizing the magnitude of persistence and variability of information over time.

The approach proposed here stems from the analysis of text similarities, which is used to compare documents in large text corpora, in order to assess the repetitions of patterns. In this respect, similarity is thus considered a measurable property assessing the degree of relation between two or more information artifacts.

There is a plethora of similarity measures (for an extensive review see [12]) to evaluate this feature; in this study we will take into account two particular vector-based metrics that have been selected for their particular properties.

A measure often used for comparing different documents is represented by the Jaccard coefficient (J), a distance metric defined as follows: given two documents A and B , let a and b the sets of terms occurring in A and B respectively. Define I as the intersection of a and b , and K as their union. Then the Jaccard similarity is the number of elements (cardinality) of I divided by the cardinality of K , thus

$$J = |I| / |K|. \quad (1)$$

Conversely, the Cosine similarity (C) is computed in the following way: let A_s and B_s be sets of terms occurring in A and B , as in the previous measure; define K as the union of A_s and B_s , and let k_i be the i -th element in K . Then the vector terms in A and B are:

$$\begin{aligned} a &= [nA(k_1), nA(k_2), \dots, nA(k_n)]; \\ b &= [nB(k_1), nB(k_2), \dots, nB(k_n)], \end{aligned} \quad (2)$$

where $nA(k_i)$ is the number of occurrences of term k_i in A , and $nB(k_i)$ is the same for B . In this respect, Cosine similarity between two original document sets is defined as

$$C = (a \cdot b) / \|a\| \|b\|, \quad (3)$$

that is the ratio between the scalar product of vectors a and b and their Euclidean norm.

Both measures are widely used for comparisons among documents on broad masses of text corpora, and broadly applied in information retrieval applications and tools [16, 5].

There is a significant difference between the Jaccard coefficient and the Cosine index: both enable to compare two texts, but the former gives more importance to non-overlapped information in the denominator: for this reason, the Jaccard index is much more sensitive to the relative size of documents. Thus, this metric allows to address size differences among documents, while the Cosine metric accounts for the preservation of the same sets of information in different texts. Consequently, when comparing changing versions of the same document in a time domain, the former metric can be used to measure the magnitude of total change, taking into account both the inclusion of new tokens/lemmas and increase/decrease in text size, while the latter signals how much of the former information has been preserved between two versions of a document.

Technically, while the Flesch readability score can be computed directly over the unwikified version of the article (a version cleaned from all wiki tags), further preprocessing of an article is needed before computing the Jaccard and Cosine coefficients: in particular the text needs to be cleaned from stop-words (such as, for instance, propositions and conjunctions) and all remaining words are then truncated into stems.

3.2 Trends

On the basis of the archive, it is possible to reconstruct the history of Simple Wikipedia. For each article in the archive we know when it was first introduced and for every month that Simple Wikipedia existed, we can find the versions of the articles in the archive that existed in that month and we can count their number and properties like their overall readability. Figure 1 in the appendix shows the result of such a reconstruction for Simple Wikipedia from January 2003 until December 2007. The figure shows an upward sloping line indicating the total number of articles in the collection for each month and a downward sloping line indicating the overall readability of the articles in the collection for each month. Note that the continuous growth in the number of articles in Simple Wikipedia is at least partially due to the fact that articles that have been removed from the collection and are not currently available anymore, do not appear in the most recent archive either. Even so, the growth in the number of articles over time is impressive. This rate of growth may be

a factor explaining why the readability indicated by the Flesch readability score shows a gradual decline: as the size of articles in Simple Wikipedia grew, it became more unwieldy and editors faced an ever harder task to maintain the standards of readability, is one interpretation that suggests itself. Looking closely, we can distinguish a phase of substantial decline in readability in the first half of 2004 followed by a more stable phase in the second half of 2004. It was in the second half of 2004 that the practice of tagging articles with the label “unsimple” first appeared. It might be that the slow down in the decline in readability could be attributed to this. With a readability index of over 70, Simple Wikipedia still scores very well given that standard English falls between 60 and 70. Still, the fact that readability continues to decline should be taken seriously into account in the close future.

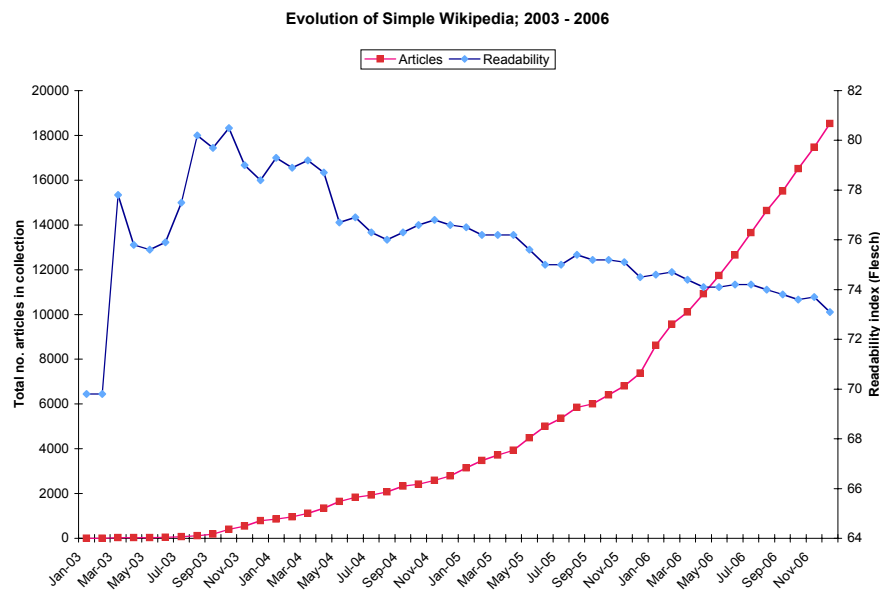


Fig. 1. Performance at Project Level – Simple Wikipedia: total number of articles per month and mean readability index per month.

Figure 2 shows how the similarity indexes change as the articles are revised. Similarity values are averaged for the various articles according to the same revision. Similarity of the *n*-th revision is computed, respectively, with reference to the first revision of the article (left side) and with respect to the previous revision (right side). Note here and for all subsequent graphical representations of similarity that the metrics are computed over the subset of all the articles which have been labeled as “unsimple” at least once and that the effects of vandal editing have been filtered out in order to improve readability.

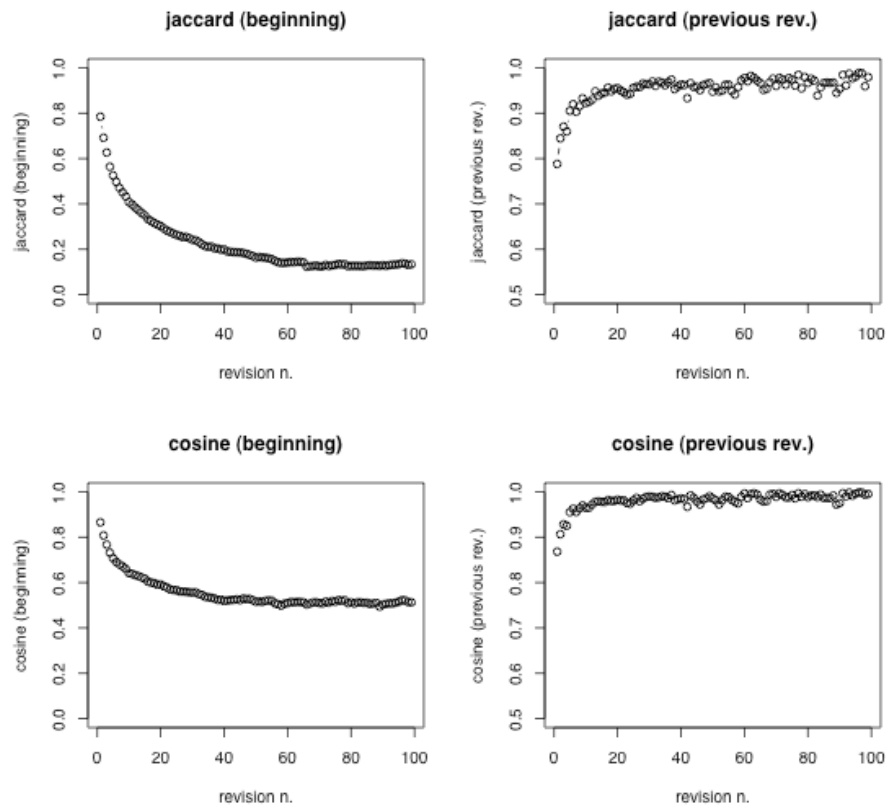


Fig. 2. Evolution of Cosine and Jaccard similarity indexes (vandalisms are filtered out).

The left side of Figure 2 shows that articles on average change considerably with respect to their initial stub and tend to converge to a stable similarity index after circa 40 revisions. Cosine similarity is higher than the Jaccard index, meaning that on average there is a relatively high persistence of initial words/lemmas compared to the high dissimilarity that is introduced overtime by introducing new words/lemmas and by increasing the text size.² The right side of Figure 2 confirms that, as the number of revisions increases, changes become more and more incremental, but since the similarity indexes oscillate at some degree near the lower bound of 1, the revision process still exhibits some, albeit minimal, degree of dissimilarity.³

² This claim takes into account that most articles start as short stub and are enlarged later on, but similar considerations holds also for those articles which are initially copied almost verbatim from en.wikipedia and that are later on simplified and shortened to comply with the Simple Wiki policies.

³ It is maybe worth to note that findings on the left side of the figure are not in contrast with right side ones if one considers that some revisions are circular (as in the frequent case of edit

Figure 3 tries to shed a light on who has been editing articles in Simple Wikipedia from 2004 until 2006. Most dominantly, it shows the growing importance of bots – indicating that automated and semi-automated scripts are relied upon more and more to keep the articles consistent. At the same time, the proportion of edits done by people who do not identify themselves to the system decreases and the proportion of edits do by people who are registered increases relative to them. Whatever causes the gradual decline in readability, it is probably not the hordes of outsiders in the case of Simple Wikipedia.

A final issue is the number of articles that have been identified as “unsimple”. This turns out to be a relatively small number: Less than 200 between 2004 and 2006. We do find however that there is a clear improvement in readability on average following the appearance of the label. Another interesting finding is that the act of tagging seems to be carried out by users with a special administrator status most of the time and hardly by anonymous outsiders (cfr. [3]).

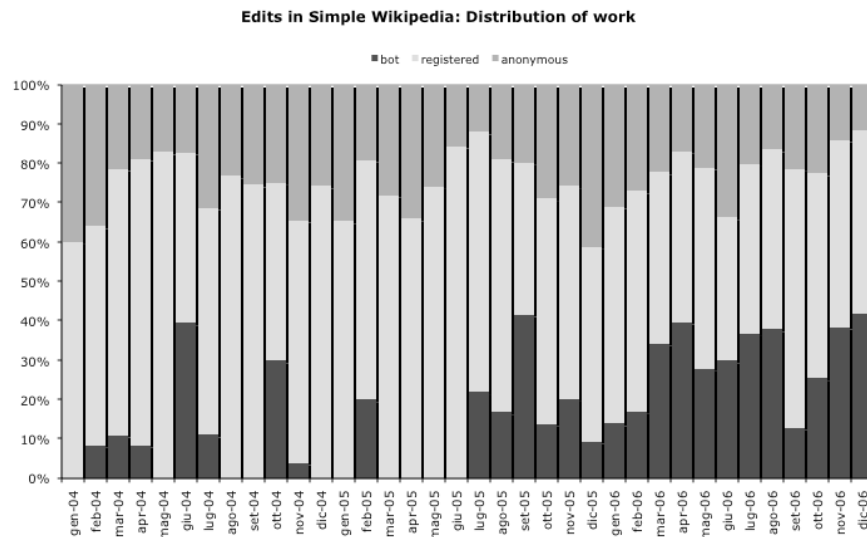



Fig. 3. Variation in editor background over time.

3.3 Triage

Despite the overall trend towards a decline in readability and the limited number of articles that have been labeled with “unsimple”, applying such a label does seem to have the desired effect in several cases. Take for instance the article on propaganda.

warring) and that in the other cases two subsequent revisions, while still presenting some degree of dissimilarity between them, might be relatively equally dissimilar to the initial revision of the article.

Figure 4 depicts the evolution of the readability over time of this page. Readability is based on the Flesch reading easy score and time is given as the number of days relative to the first occurrence of the tag “unsimple” on the page. In addition, the dotted vertical lines indicate the period during which the page was labeled with the tag “unsimple” and the shapes of the points correspond to the different contributors to the page.



Randywombat. In the year that followed we see edits by five edits by non-identified users and two by identified users. There are also three edits by bots, but these do nothing more than adding so-called inter-wiki links to similar pages in other languages. In November 2005, the user Heroismic, a high-school junior (according to his/her user-page) who has not edited the page before, comes along and labels the article “unsimple”. In the following months, we see a concerted effort by the users Eptalon, according to his/her user-page a Simple Wiki administrator from Central Europe, and AmanitaMuscaria to improve the page. On 22 February 2006 Eptalon decides that the page has become readable enough and removes the label “unsimple”.

Figure 5 shows the development of readability for few more selected articles. In this figure, the readability of the articles is represented by a ‘o’ while the article has the label “unsimple” and a ‘+’ otherwise. Sometimes readability increases but no one bothers to remove the “unsimple” label. In other cases, the label leads to some improvement in readability initially but contributors fall back on bad habits later. And, sometimes, the label seems to be ignored altogether. The interesting question here is not just how likely it is that this kind of “treatment” by labeling succeeds sometimes, but more why it succeeds in some cases and not in others (for more on this topic see also below in the final section).

Similarly to what we did for readability, we can apply the similarity metrics to selected articles in order to single out different editing processes at work in “simple” vs. “unsimple” regimes. Figure 6 shows the similarity indexes for the article on the year 2001 (“2001”): here for any given revision similarity is computed with respect to the first revision of the ongoing regime. Both Cosine and Jaccard similarity indexes drop abruptly at the beginning; this is largely due to the fact that the article starts as a very short stub on 5/7/04 and it is subsequently substituted at its third revision (on 5/15/04) with a verbatim copy from en.wikipedia. The article remains in the “simple” regime for a total of 37 revisions in the period 5/7/04 – 5/7/06, at the end of which $C=0.32166$ and $J=0.00179$. The unsimple tag is placed by an admin and remains in its place for the subsequent 22 revisions until the article status switches back to “simple” on 12/4/06 for all the remaining available observations (21 revisions up to 6/11/07). If one look closely to the figure, there is a considerable difference in end of regime measures of similarity. At the end of the unsimple regime $C=0.97457$ and $J=0.92732$, while at the end of the second simple regime $C=0.62334$ and $J=0.88344$.

These large differences seem to suggest the existence of rather different patterns of article revision during the “unsimple” regime compared to the “simple” one. In particular, what seems to be particularly counterintuitive is the realizing that the troublesome “unsimple” regime seems to be characterized by a sensibly slower pace of change. While this finding doesn’t hold in principle for all the articles (and calls for more accurate analysis at the single article level), it seems to be confirmed, on average, at the aggregate level.

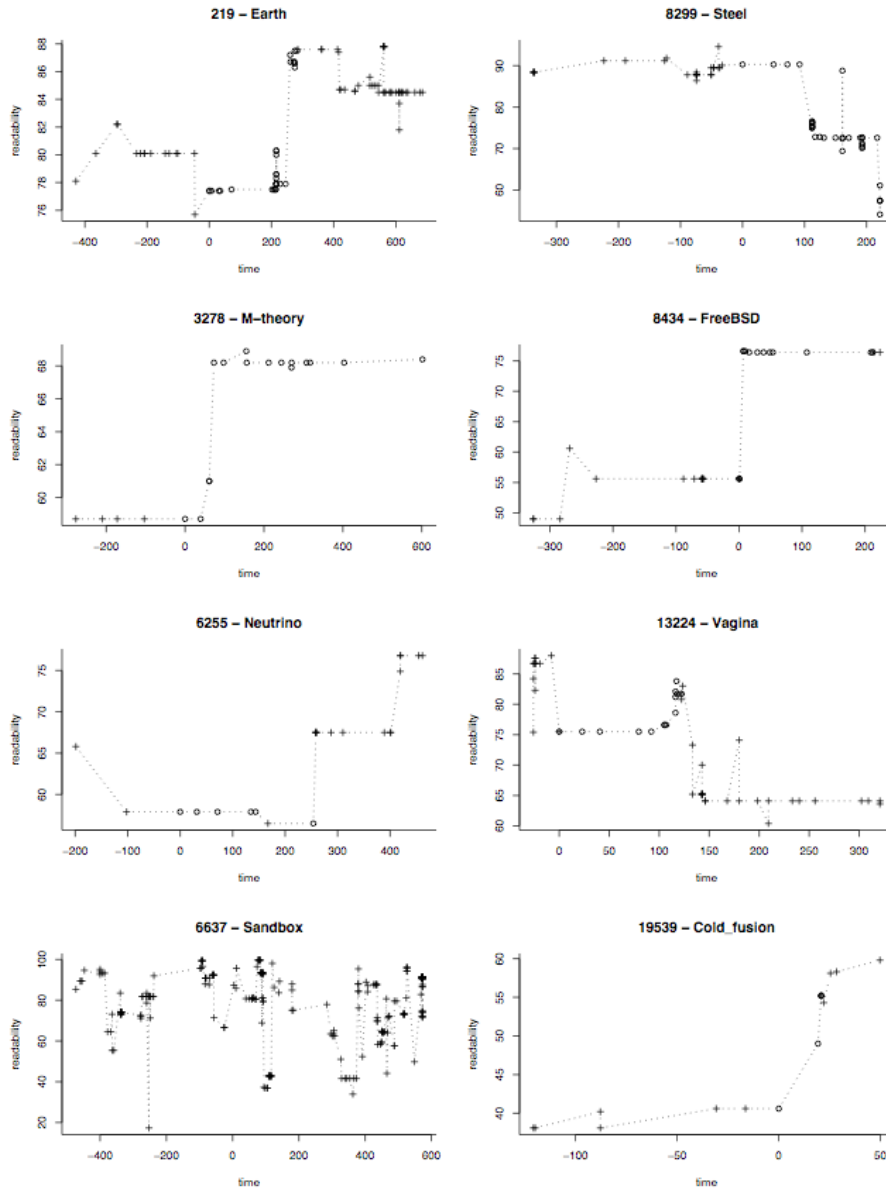


Fig. 5. Development of Readability in Selected Articles.

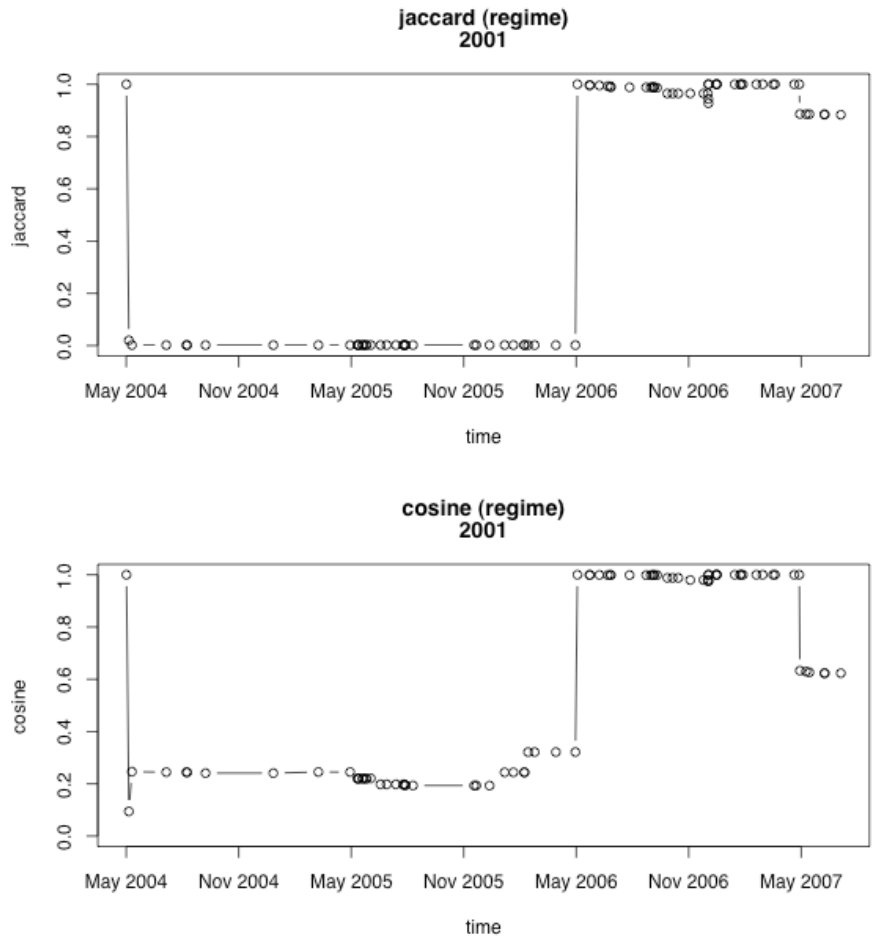


Fig. 6. Treatment of "2001" in Simple Wikipedia: similarity indexes in the simple/unsimple regimes (metrics are computed for any revision with respect to the first revision of the outstanding regime).

Figure 7 shows the distributions of C and J indexes (computed according to the method described for Figure 6) for the first 3 regimes., where regime 1 corresponds to the initial period of editing, regime 2 starts at the first time the article is tagged as "unsimple" and regime 3 starts when the tag is removed (and lasts, eventually, until the "unsimple" tag is placed again). The figure seems clearly to suggest the existence of a two-speed process according to which in the first "simple" regime dissimilarity increase at a very fast pace (mean values: $J=0.549$, $C=0.726$) while revisions in the "unsimple" regime 2 are more incremental (mean values: $J=0.764$, $C=0.9$), and dissimilarity increases again in the subsequent "simple" regime 3 (mean values: $J=0.703$, $C=0.834$). All these differences between two subsequent regimes are

statistically significant (Wilcoxon signed rank test, one sided, $\alpha=5\%$ or lower). Overall, this finding suggests that during non-pathological periods, the editing style seems to favor large revisions making the article progressively more complicated. On the contrary, when the article is tagged as “unsimple”, efforts towards regaining simplicity goes into the direction of more careful and incremental revisions.

We claim that these findings call for further inspection, in particular we find promising to cross readability and similarity analysis looking for a more articulated explanation of the above mentioned differences between “simple” and “unsimple” regimes.

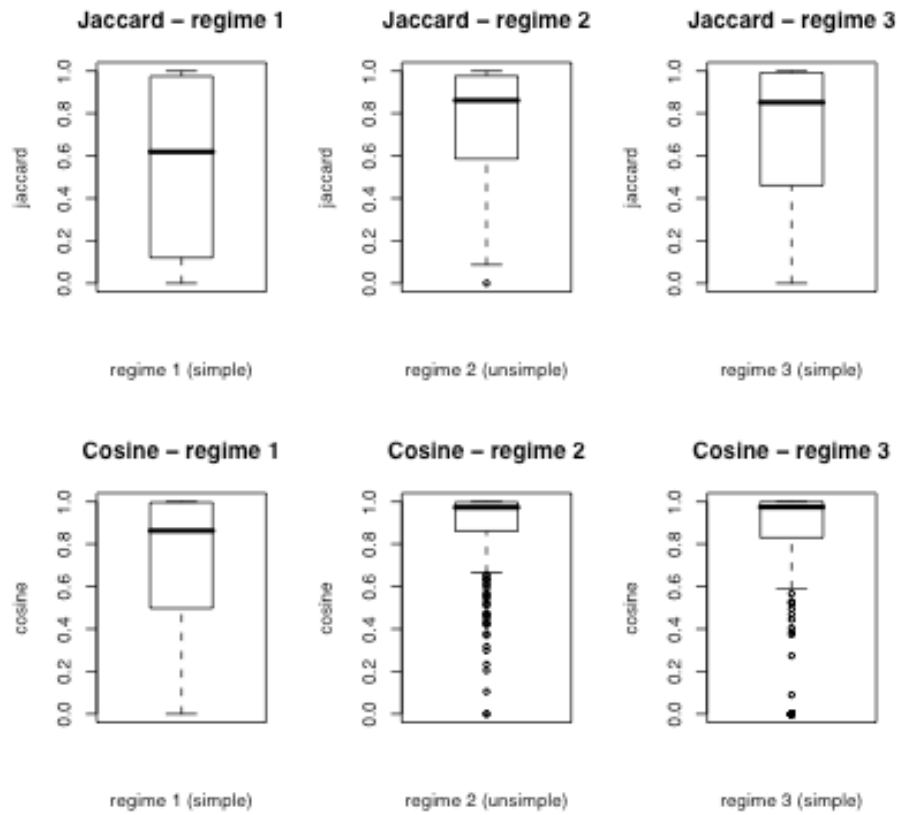


Fig. 7. Box-and-whisker plots of J and C similarity in “unsimple” and “simple” regimes in Simple Wikipedia.

4 Conclusions

The quality and sustainability of Wikipedia is a topic in which scholars of all kinds are getting more and more interested. In this paper, we have provided a review of

recent studies of Wikipedia and we have proposed a new approach to study its performance that takes advantage of the way in which this encyclopedia project is organized. In particular, we propose to look for signs within the Wikipedia archive that specific shortcomings of articles are identified and then assess how well these shortcomings are addressed. We have illustrated our approach with a pilot study of Simple Wikipedia, a spin-off of the main Wikipedia that focuses on articles that are easy to read. We found that while the editors of this collection have a difficult time in keeping articles simple, the label “unsimple” that they employ to single out articles that are particularly problematic still provides a good entry-point in studying how this particular community manages to keep its encyclopaedia readable.

With the Simple Wikipedia study as a “proof of concept”, the time is now ripe to move to other, larger, Wikipedia projects and extend the study to indicators of shortcomings like “controversial” and “stub”. In particular, we find promising to carry out a survival analysis on several variables describing the articles’ state at the point of triage and subsequent edits during the treatment to study the effectiveness of labeling as an organizational practice.

References

1. A. Ciffolilli, Phantom authority, self-selective recruitment and retention of members in virtual communities: The case of Wikipedia, *First Monday*, Vol. 8 (2003) No. 12.
2. J. Čuhalev, Ranking of Wikipedia articles on search engines for searches about its own articles, 2007, <http://www.kiberpipa.org/~gandalf/blog-files/wikistatus/wikistatus.pdf> (last retrieval November 30, 2007).
3. M. den Besten, J.-M. Dalle, Collaborative maintenance of collections – keep it simple: a companion for Simple Wikipedia?, *Industry and Innovation*, Vol. 15 (2008) No.2, pp.169–178.
4. P. Denning, J. Horning, D. Parnas, L. Weinstein, Wikipedia risks, *Communications of the ACM*, Vol. 48 (2005) No. 12, pp. 152–152.
5. I.S. Dhillon, D. Modha, Concept decompositions for large sparse text data using clustering, *Machine Learning*, Vol. 42 (2001) No. 1, pp. 143–175.
6. P. Duguid, Limits of self-organization: Peer production and laws of quality, *First Monday*, Vol. 11(2006) No. 10.
7. S. Finkelstein, Read me first: Inside, Wikipedia is more like a sweatshop than Santa’s workshop, *The Guardian*, December 6, 2007.
8. R. Flesch, *How to Write Plain English*, Harper and Row, New York, 1979.
9. J. Giles, Internet encyclopaedias go head to head, *Nature*, Vol. 438 (2005) No. 7070, pp. 900–901.

10. A. Kittur, B. Suh, B.A. Pendleton, E.H. Chi, He says, she says: conflict and coordination in Wikipedia, CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems, New York, ACM Press, 2007, pp. 453–462.
11. J. Lanier, Digital Maoism: The Hazards of the New Online Collectivism, *Edge*, May 30, 2006.
12. L. Lee, Measures of distributional similarity, Proceedings of the 37th conference on Association for Computational Linguistics, Montreal, Association for Computational Linguistics, Morristown, 1998, pp. 25–32.
13. J. Mateos Garcia, W.E. Steinmueller, Applying the open source development model to knowledge work, INK Open Source Research Working Paper 2, SPRU – Science and Technology Policy Research, University of Sussex, 2003.
14. A. Orłowski, Farewell, Wikipedia? *The Register*, March 6, 2007.
15. J. Reagle, A Case of Mutual Aid: Wikipedia, Politeness, and Perspective Taking, 2004, <http://reagle.org/joseph/2004/agree/wikip-agree.html> (last retrieval May 10, 2007).
16. G. Salton, M. McGill, Introduction to Modern Information Retrieval, McGraw-Hill, New York, 1983.
17. F. Viégas, M. Wattenberg, K. Dave, Studying cooperation and conflict between authors with history flow visualizations, Proceedings of the SIGCHI conference on Human factors in computing systems, 2004, ACM Press, New York, pp. 575–582.
18. J. Voß, Measuring Wikipedia, in: P. Ingwersen, B. Larsen (eds): Proceedings of the 10th International Conference of the International Society for Scientometrics and Informetrics, Karolinska University Press, Stockholm, 2005.