

Wikibugs: using template messages in open content collections

Loris Gaio
University of Trento
Department of Computer and
Management Science
Via Inama 5
I-38100, Trento, Italy
+39 0461 882142
lgaio@cs.unitn.it

Alessandro Rossi
University of Trento
Department of Computer and
Management Science
Via Inama 5
I-38100, Trento, Italy
+39 0461 882101
arossi@cs.unitn.it

Matthijs den Besten
Ecole Polytechnique
Centre de Recherche en Gestion
ENSTA PREG CRG 32
boulevard Victor
75739 Paris Cedex 15, France
+33 1 45 52 64 27
[matthijs.den-
besten@polytechnique.edu](mailto:matthijs.den-besten@polytechnique.edu)

Jean-Michel Dalle
Université Pierre et Marie Curie
4 place Jussieu
75252 Paris cedex 05
France
+33 1 44 18 07 15
jean-michel.dalle@upmc.fr

ABSTRACT

In the paper we investigate an organizational practice meant to increase the quality of commons-based peer production: the use of template messages in wiki-collections to highlight editorial bugs and call for intervention. In the context of SimpleWiki, an online encyclopedia of the Wikipedia family, we focus on {complex}, a template which is used to flag articles disregarding the overall goals of simplicity and readability. We characterize how this template is placed on and removed from articles and we use survival analysis to study the emergence and successful treatment of these bugs in the collection.

Categories and Subject Descriptors

K.6.1 [Management of Computing and Information Systems]: Project and People Management – *management techniques, systems development.*

General Terms

Management.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WikiSym '09, October 25-27, 2009, Orlando, Florida, U.S.A.
Copyright © 2009 ACM 978-1-60558-730-1/09/10...\$10.00.

Keywords

Commons based peer production, Wikipedia, wiki, survival analysis, quality, bug fixing, template messages, coordination.

1. INTRODUCTION

Issues pertaining to the reliability of open content collections are at the core of the agenda of both scholars and practitioners interested in commons-based peer production. As put forward by Larry Sanger, Wikipedia co-founder:

“It’s fun, first of all. But it can be fun for intellectually serious people only if we know that we’re creating something of quality. And how do we know that? The basic outlines of the answer ought to be fairly obvious to anyone who has read Eric S. Raymond’s famous essay on the open source movement, ‘The Cathedral and the Bazaar’. Remember, if we can edit any page, then we can edit *each other’s work*. Given enough eyeballs, all errors are shallow. We catch each other’s mistakes and enjoy correcting them.” [19]

Others have been more agnostic regarding the possibility of large mass peer screening to act as a substitute for source authoritativeness as a means for assuring quality [13]. Obviously, as far as trustworthiness is concerned, content peer production has also its share of skeptics in the scientific literature [10], in practitioners’ view [14] and in popular media [5]. Despite some exceptions [12], this lively debate has mostly been fueled by claims that have still to move towards the stage of sound empirical validation.

We build from previous empirical research in the field that has started to shed light on the role of institutions and organizational practices in channeling the largely unstructured efforts of voluntary contributors [8-9-16]. According to this line of research, peer production within wiki platforms makes extensive use of template messages – standard info-boxes placed on top of a given page – as coordination tool which ease the contribution to the production process of the various participants. In Wikipedia, for instance, there is an overwhelming number of templates, a.k.a. tags, which are used as a means to facilitate various goals and activities, such as to flag particular anomalies and dysfunctions of pages (e.g., violations of common policies or guidelines), and to call for specific actions for contributors (e.g., cleaning up, improvements in the organization of the text, and so on).

In what follows, we concentrate on templates signaling breaches of important policies or guidelines – consensual standards and advisory statements which every editor should bear in mind when editing an article in the collection – and, drawing a parallel with open source software methodologies, we treat placing/removing such templates from the text of an article a process similar to filing /closing a *bug* report in software development. Indeed, the use of templates signals the existence of a *problem* associated with the page that it flags, by highlighting this problem in a very similar way to the reporting of a bug in software contexts. While placing a template on a page is similar to reporting a bug, applying modifications in such a way to remove the problem is similar to bug fixing. We therefore suggest to denote problems associated with the content developed using wikis as *wikibugs*, not only in order to stress similarities with bugs and to build upon existing knowledge associated with software bug management, but also to stress the fact that there is a more general issue about problem-solving in online communities, which both software bugs and wikibugs precisely instantiate in two different contexts. Focusing on bug fixing in open source development, recent research [4-6-7] has indeed shed light on various important organizational issues, e.g., notably about the drivers of coders' attention, which accounts for the more or less quick fixing of bugs. Conversely, the same topic has received much less coverage in the realm of open content peer production.

But we are also perfectly conscious that wikibugs are not bugs, and that the organizational processes and practices according to which they are reported, dealt with and eventually solved are different: on the very contrary, one of the outcomes that we expect from studying wikibugs has to do with analyzing more precisely the specificities of (distributed) problem solving in the context of wikis, and with comparing it with other forms of distributed problem solving in other online communities. Furthermore, we believe that the analysis of wikibug fixing activities is crucial in improving our understanding how and to what extent it is possible to reconcile the apparent contrast between spontaneous collaborative authorship and quality assurance of a wiki collection.

In particular, in this paper we aim at understanding through descriptive and survival analysis which are the variables in an article production process which account for the emergence of a wikibug or have influence on how the wikibug is fixed. In particular, we model the dynamic of tagging as a survival process, linking the probability of entry/exit of a page into the “pathological state” to various explanatory variables. These variables can be traced back to three large families: measures of effort and intensity of work, measures of participation of

users/division of labor and measures related to morphological features of pages (size, readability, similarity overtime, and so on). According to this framework, we perform survival analysis (i) on the duration of the initial non-pathological state, thus studying which variables shorten or increase the amount of time needed for an encyclopedic article to develop undesirable features (ii) on the duration of the pathological state, exploring how different variables affects the treatment or the persistence of such undesirable features.

This paper is organized as follows: we first offer some justification for the choice Simple Wikipedia as the empirical field of investigation. We then describe the methods used to retrieve data and how we performed the analysis. Then we present our main findings, distinguishing between descriptive and survival analysis. Finally we offer some concluding remarks and suggestions for further research on related topics.

2. DATA

In order to perform our study we mined data from Simple Wikipedia, an online encyclopedia, which belongs to the larger Wikipedia project and is intended for readers whose first language is not English.

The fundamental reason of the choice of Simple Wikipedia over various other publicly available wiki-based collections lies in the strong commitment by the active participants in the activities of collaborative editing to a writing style which poses a strong emphasis on simplicity and readability [24-25]. Accordingly, the template `{{complex}}` – which in the early days of the collection was labeled `{{unsimple}}` – is used by editors in order to signal that a particular article is unsatisfactory as far as readability is concerned. This is certainly not the only instance of use of template messages that closely parallels how bug reports work in software development (e.g., a prominent Wikipedia template in this respect is represented `{{NPOV}}`, which flags the breach of the fundamental policy of writing articles according to a neutral point of view). Nevertheless, an additional rationale pushing further the argument for choosing `{{complex}}` from Simple Wikipedia over other alternatives is represented by the possibility of computing measures of simplicity/readability, derived from the computational linguistic tradition, which can be used as objective appraisal of the gravity of a bug and of the effectiveness of the work done to fix it.

Data mining and preparation of the database for the analysis. We used the July 2008 archive of Simple Wikipedia, available at: <http://downloads.wikimedia.org>, which for every revision made on an article page lists the following data: the user-id of the editor (IP address in case of anonymous edits), date and time of the edit, comments made by the editor and the full (wiki markup) text of the revision. We selected from the archive all the revisions corresponding to article pages which had been tagged at least once with the `{{unsimple}}`/`{{complex}}` templates. In order to avoid biases due to very short series for some datapoints in the survival analysis, we restricted the analysis to article pages which had been revised at least 15 times. For each article page we limited our extraction to all revisions belonging to the interval spanning from the first edit to the revision antecedent to the removal of the `{{complex}}` template¹. After this selection, we ended up with

¹ In the case of pages in which the `{{complex}}` tag has never been removed (a.k.a. censored pages) we took all the available revisions.

378 article pages for the analysis. De-wikification of the text and categorization of registered users (in terms of administrators, bots and “plain” registered users) have been performed according to previous literature [8-9]. While some studies on the English Wikipedia have shown that actual changes in a given article page are sometime the result of longer discussions occurring at the level of the corresponding talk page [15, 23], the use of talk pages as a means to anticipate and discuss actual changes is relatively infrequent in Simple Wikipedia. Thus, the dataset employed in our analysis is restricted to data from article pages only.

Vandal filtering. Previous work has highlighted the short life span for vandal edits in wiki-collections [22]. While it is generally reassuring to know that there is only a limited impact of these malicious activities on the quality of the whole archive, at the same time when studying the process of development of articles one still has to carefully evaluate whether vandal edits might introduce distortions in the interpretation of the data. In our case filtering out vandal edits was essential in order to be able to perform unbiased measures of work activity related to article pages (e.g., number of revisions, number of unique contributors, etc.). While we were aware of the existence of algorithms for the automatic detection of vandalism [17], given the relatively limited number of revision involved, we preferred to manually check for vandalisms. We performed this activity both using comment analysis (in order to single out reverts which were explicitly accounted by editors as fixes to vandal edits) and MD5 hash (computed over the full text of a revision) comparisons across subsequent revisions of an article page. Overall, we filtered out from the dataset 11% of revisions which were categorized as vandalisms or revert edits fixing previous vandalizations.

3. METHODOLOGY

We applied survival analysis to study two different albeit intertwined phenomena: (i) transition of article pages from the initial “simple” phase (from now on: regime 1) to the subsequent “unsimple” phase (from now on: regime 2) and (ii) exit from regime 2. The observation periods are, respectively, from the very first version of an article page to the revision in which the template `{{complex}}` appears, and from the latter to the revision in which the template is edited out. By definition of the sample, for the first event (exit from regime 1) all observations are uncensored, while for the second event some observations are censored, meaning that in some cases the template has never been removed from the article page.

In order to perform the survival analysis we used the dataset extracted in the previous Section to calculate the following variables:

- *duration of regimes:*
 - `duration1`, `duration2`: duration (days) of regime 1, 2;
 - `reactime`: reaction time (days) measured as time between tagging and the first subsequent revision of the article;
- *intensity of efforts and division of labor:*
 - `revs1`, `revs2`: number of edits in regime 1, 2;

Also, we did not consider instances of repeated flagging of one page, where one page, after returning in the “simple” regime, is flagged once again as complex.

- `regrevs1`, `regrevs2`: number of edits by registered users in regime 1, 2;
- `admrevs1`, `admrevs2`: number of edits by administrators in regime 1, 2;
- `anonrevs1`, `anonrevs2`: number of edits by anonymous users in regime 1, 2;
- `botrevs1`, `botrevs2`: number of edits by bots in regime 1, 2;
- `anon1`, `anon2`: number of unique anonymous editing in regime 1, 2;
- `admin1`, `admin2`: number of unique administrators editing in regime 1, 2;
- `bot1`, `bot2`: number of unique bot editing in regime 1, 2;
- `reg1`, `reg2`: number of unique registered users editing in regime 1, 2;
- *textual measures:*
 - `char1`, `char2`: character count at the end of regime 1, 2;
 - `word1`, `word2`: word count at the end of regime 1, 2;
 - `read1`, `read2`: Flesch readability score at the end of regime 1, 2;
 - `cosim`, `jaccsim`: Cosine and Jaccard similarity measures computed between the end of regime 1 and the end of regime 2.

Finally, readability was computed via the Flesch Readability Index and we used the Cosine and the Jaccard Indexes as similarity metrics (for the details the reader is referred to [9]).

4. RESULTS

4.1 Descriptive statistics

Table 1. Summary statistics for the `{{complex}}` articles

	Min.	1 st Qu.	Median	Mean	3 rd Qu.	Max
duration (days)	273.7	823.5	1082	1125	1480	2481
revisions	5	16	29	53.73	59	559
editors	2	10	17	30	33.5	222

Table 2. Summary statistics for regime 1 and regime 2 durations (days)

	Min.	1 st Qu.	Median	Mean	3 rd Qu.	Max
R1	0.0003	3	260	356	622	1765
R2 (all)	0.0023	66	269	356	625	1401
R2 (closed only)	0.0023	24	122	178	270	962
R2 (ongoing only)	273	614	653	727	811	1401

Table 1 presents some descriptive statistics on the duration, number of revisions and number of editors computed over the complete lifetime of the 378 articles tagged at least once with the `{{unsimple|complex}}` template. Table 2 compares durations of regime 1 and regime 2. The comparison shows similar distributions for the right side of regime 1 and regime 2, while the

first quartile of regime 1 shows that a considerable share of articles is tagged just right after their inception. For regime 2, Table 2 also distinguishes between uncensored articles ($n=255$, for which the template has been removed at some point) and censored articles ($n=123$, still tagged as complex at the time of the dataset collection). This comparison shows that the uncensored group is treated in a relatively quick way, while for the censored one the duration of regime 2 is quite a bit longer.

4.2 Findings for regime 1

According to a Kaplan Meyer estimate, regime 1 seems to fit quite well to a Cox Proportionality Hazard class model. Figure 1 depicts the survival function, using Kaplan Meyer estimation, fitted using duration times for regime 1.




Figure 1. Survival function for regime 1 (Kaplan Meyer estimator).

In order to assess the different effects of covariates in the termination of regime 1, we start considering the impact of division of labor, and in particular the incidence of efforts by different kind of users towards duration of regime 1. For this purpose we need preliminary to screen for the possible existence of multicollinearity issues between the various variables.

Table 3 summarizes correlations existing between the duration of regime 1, and two families of covariates: variables related to effort (edits) exerted by different categories of participants and variables measuring the number of participants (for such categories). The table shows the existence of a strong correlation between participants and edits for each category considered. Hence, the use of both families together in the survival estimation should be avoided.

Consequently, a CoxPH model has been fitted in order to explain the impact of the three families of covariates depicted in Section 1. Overall, only the variables pertaining to intensity and division of labor seem to have a significant effect in explaining the length of regime 1, while variables regarding other features of the pages, such as size, readability, similarity have no explanatory power. For sake of brevity we present the final models only, which are summarized in Table 4.

Table 3. Correlation matrix for regime 1 covariates

	duration1	regrevs1	admrevs1	anonrevs1	botrevs1	reg1	adm1	anon1	bot1
duration1	1	.401	.398	.511	.615	.710	.575	.618	.640
regrevs1	.401	1	.296	.459	.304	.651	.364	.465	.324
admrevs1	.398	.296	1	.525	.387	.467	.807	.590	.367
anonrevs1	.511	.459	.525	1	.447	.652	.605	.912	.406
botrevs1	.615	.304	.387	.447	1	.525	.573	.506	.895
reg1	.710	.651	.467	.652	.525	1	.588	.755	.531
adm1	.575	.364	.807	.605	.573	.588	1	.673	.537
anon1	.618	.465	.590	.912	.506	.755	.673	1	.441
bot1	.640	.324	.367	.406	.895	.531	.537	.441	1

In Model 1 the duration of regime 1 is negatively affected by the number of revisions by all categories of contributors. Similarly, there is a negative impact on duration when considering the number of different contributors per category (Model 2). The latter model seems to have a higher descriptive power as far as Rsquare and model tests are concerned.

Table 4. Survival Analysis on Regime 2 Inception

Variable	Model 1	Model 2
regrevs1	-0.039*** (0.0129)	-
admrevs1	-0.037* (0.0203)	-
anonrevs1	-0.027** (0.0089)	-
botrevs1	-0.105*** (0.0150)	-
reg1	-	-0.147*** (0.0267)
adm1	-	-0.135*** (0.0510)
anon1	-	-0.025* (0.0157)
bot1	-	-0.254*** (0.0375)
Rsquare	0.370	0.492
L ratio	175	256
Wald	122	177
Logrank	120	185

p-values significance: * <0.1 , ** <0.05 , *** <0.01

Overall the two models seem to suggest that both the level of effort on a page (in terms of revisions) and the number of participants in the editing process anticipate the emergence of readability concerns. At this point of the analysis it is still difficult to judge whether this shortening is more due to a variant of the Linus' law (more eyeballs resulting in the anticipatory detection of a defect) or rather due to diminishing returns related with increases in the number of contributors. While the second model seems to be more ambiguous in this respect, the first one seems more clearly to suggest a connection between increases in work intensities and the emergence of a bug as the result of coordination conflicts. Nevertheless this issue seems to be worth of further scrutiny.

4.3 Findings for regime 2

Similarly to the previous regime, for regime 2 durations a Kaplan Meyer estimate has been computed (Figure 2) and the model seem again to fit quite well a Cox Proportionality Hazard class model.




Figure 2. Survival function for regime 2 (Kaplan Meyer estimator)

Table 5 confirms the existence of a strong correlation between the number of revisions made by different classes of participants and the number of participants (for the same classes), again suggesting to avoid the use of both families in the same model estimation in order to avoid for multicollinearity problems.

Similarly to what has been done for regime 1, we test for the same hypotheses related to efforts and division of labor; we look whether the total number of revisions and the number of different contributors in the various classes do play a significant role in exiting from regime 2.

Results are reported in Table 6. Here the variables related to the intensity of efforts (number of revisions) are not significant with the exception of revision made by bots (Model 1). On the contrary, all classes of users are significant when considering the number of different contributors per category (Model 2). In particular the shortening of the pathological regime seems to be affected by the presence of administrators, registered and bot

users, while the presence of anonymous users seems to delay the fixing process.

Table 5. Correlation matrix for regime 2 covariates

	duration2	duration1	regrevs2	admrevs2	anonrevs2	botrevs2	reg2	adm2	anon2	bot2	react
duration2	1	.000	.307	.376	.327	.599	.438	.543	.355	.706	.199
duration1	.000	1	.182	.190	.197	.178	.220	.226	.204	.159	-.066
regrevs2	.307	.182	1	.816	.792	.467	.887	.710	.803	.358	-.031
admrevs2	.376	.190	.816	1	.929	.618	.904	.897	.931	.490	-.038
anonrevs2	.327	.197	.792	.929	1	.571	.888	.842	.988	.436	-.038
botrevs2	.599	.178	.467	.618	.571	1	.635	.645	.583	.922	-.061
reg2	.438	.220	.887	.904	.888	.635	1	.849	.908	.525	-.028
adm2	.543	.226	.710	.897	.842	.645	.849	1	.866	.578	.006
anon2	.355	.204	.803	.931	.988	.583	.908	.866	1	.453	-.028
bot2	.706	.159	.358	.490	.436	.922	.525	.578	.453	1	-.034
react	.199	-.067	-.031	-.038	-.038	-.061	-.028	.006	-.028	-.034	1

Table 6. Survival Analysis on Regime 2 Termination

Variable	Model 1	Model 2	Model 3
regrevs2	-0.014 (0.0098)	-	-
admrevs2	0.015 (0.0196)	-	-
anonrevs2	0.007 (0.0064)	-	-
botrevs2	-0.091*** (0.0103)	-	-
reg2	-	-0.060* (0.0373)	-0.084** (0.0379)
adm2	-	-0.233*** (0.0516)	-0.262*** (0.0519)
anon2	-	0.074*** (0.0131)	0.081*** (0.0135)
bot2	-	-0.224*** (0.0222)	-0.228*** (0.0220)
duration1	-	-	.0004** (0.0001)
Rsquare	0.320	0.499	0.537
L ratio	146	262	291
Wald	83.4	152	174
logrank	89	189	213

Similarly to the previous Subsection, other covariates (in particular the textual-related covariate) have no incidence on the survival process. In particular, the reaction time to flagging has a negligible impact on the duration of regime 2 (as we found in related research not presented here). Model 3 allows to introduce in the survival the duration of regime 1 (that can be also thought as the overall life of the page at starting of regime 2) as a covariate. This variable is significant and affects positively the duration of regime 2. A possible interpretation is that the older the page at time of flagging, the more difficult is to solve successfully readability issues.

A final remark that is worth making is on the variable measuring the efforts made by users which originally tagged the page. This variable is not significant, thus hinting to a quite different story with respect to open source development as far as to bug fixing is concerned, and reinforcing a view of open content creation communities as made more by “passers-by” users, rather than by contributors which commit themselves to a particular artifact on a long term perspective.

5. DISCUSSION

The importance of “passers-by” was first remarked in the controversial paper by Antony *et al.* [1] who argue that high quality contributions often come from anonymous “Samaritans” who contribute infrequently. Their analysis focused on the addition of new content. On the basis of the findings of Bryant *et al.* [2] we would have expected to find that a lot of the subsequent cleaning is carried out by regular contributors and especially by those who have acquired an administrator status. However, in our regressions we did not find evidence for this supposition. Thus, more than common identity or bonding [18], mere attention seems to be the most important determinant for the swift resolution of the wikibugs we studied.

In any case, already the fact that for many articles in our collection it takes a very long time for the defect-indicator `{{complex}}` to be removed indicates that the decentralized governance of Wikipedia described in [11] is not without its challenges. A possible way to address these challenges is to invest more in technologies, particularly those aimed at directing the attention of potential contributors [3, 8].

Then again, we feel that what is needed at this stage is most of all more analysis of the existing archive in order to get a better understanding of how Wikipedia works. For instance, there is no point in directing attention if the attention is directed towards articles which are not open to outside contributions in the first place [21]. In fact, the Wikipedia archives do indeed constitute a rich source for all kinds of investigations into human behavior (e.g., [20]). But these empirical investigations require quite a bit of effort and the results that we have been able to present so far can still only be classified as preliminary.

6. CONCLUSIONS

In the paper we studied how bugs emerge and are treated in collaborative authorship collections, and we report our preliminary findings on a survival analysis performed on the durations of entry and exit times of pages from a pathological treatment.

As far as regime 1 is concerned, we showed that entry in the pathological regime is affected both by the number of users and their efforts, and the former model seems to be relatively more

robust. Conversely, no structural feature of pages like size, readability, similarity, and so on are helpful in explaining the `{{complex}}` tagging. Overall, survival findings might highlight the existence of competing explanations regarding the shortening of the duration of regime 1 (complexity/coordination issues vs. “eyeballs” hypothesis), which call for further scrutiny.

Regarding regime 2, exit from the pathological state seems to depend on factors related on the number of participants only. In particular, while anonymous users have detrimental effects, all three categories of registered users seem to help in sorting the readability issue.

Finally, we mentioned that both entry and exit cannot be traced back neither to reaction time measures, nor other structural features of pages, such as readability, similarity, and so on. In this respect we think that, other statistical models, i.e. event analysis, might represent a more suitable way to study in a more dynamic way their effect on pages being tagged.

The research reported here can be regarded as the first step towards a more comprehensive understanding on the effectiveness of various organizational practices on bug spotting and bug fixing in collaborative authored open content collections. In this line, we find promising to apply the survival analysis framework to the study of other template messages, which signal the breach of important policies (such as above mentioned `{{NPOV}}`) and to larger datasets, such as the ones which can be derived by mining the English Wikipedia.

7. ACKNOWLEDGMENTS

We would like to thank three anonymous referees for helpful comments and Nahid Mahbub which provided help in data collection. Loris Gaio and Alessandro Rossi received partial funding from MIUR (projects FIRB03 and PRIN05).

8. REFERENCES

- [1] Antony, D., Smith, S.W. and Williamson, T. 2005. Explaining quality in Internet collective goods: zealots and good samaritans in the case of Wikipedia. Technical report, Dartmouth College, Hannover, NH.
- [2] Bryant, S.L., Forte, A. and Bruckman, A. 2005. Becoming wikipedian: transformation of participation in a collaborative online encyclopedia. In GROUP’05:6–9.
- [3] Cosley, D., Frankowski, D., Terveen, L. and Riedl, J. 2006. Using intelligent task routing and contribution review to help communities build artifacts of lasting value. In: Proceedings of the SIGCHI conference on Human Factors in computing systems:1037–1046, Montréal, Québec, Canada. ACM.
- [4] Crowston, K. and Scozzi, B. 2004. Coordination practices for bug fixing within FLOSS development teams. In: Proceedings 1st CSAC. Porto, Portugal.
- [5] Cuozzo, S. 2008. Urban Myths: What Wikipedia gets wrong about NYC. New York Post. http://www.nypost.com/seven/08242008/postopinion/opedcolumnists/urban_myths_125773.htm?page=0. Accessed 21 January 2009.
- [6] Dalle, J.-M. and den Besten, M. 2007. Different bug fixing regimes? A preliminary case for superbugs. In: Proceedings of the 3rd International Conference on Open Source Systems. Limerick, Ireland.

- [7] Dalle, J.-M., den Besten M. and Masmoudi, H. 2008. Channelling Firefox developers: mom and dad aren't happy yet. In: Open Source Development, Communities and Quality, B. Russo, E. Damiani, S. Hissam, B. Lundell, G. Succi, Eds., IFIP International Federation for Information Processing; 275. Boston, Springer:265–271.
- [8] den Besten, M. and Dalle, J.-M. 2008. Keep it Simple: A Companion for Simple Wikipedia? *Ind. Innovat.* 15(2):169–178.
- [9] den Besten, M., Rossi, A., Gaio, L., Loubser, M. and Dalle, J.-M. 2008. Mining for practices in community collections: finds from Simple Wikipedia. In: Open Source Development, Communities and Quality, B. Russo, E. Damiani, S. Hissam, B. Lundell, G. Succi, Eds., IFIP International Federation for Information Processing; 275. Boston, Springer:105–120.
- [10] Denning, P., Horning, J., Parnas, D. and Weinstein, L. 2005. Wikipedia risks. *CACM* 12(48):152.
- [11] Forte, A. and Bruckman, A. 2008. Scaling consensus: increasing decentralization in Wikipedia governance. In: Proceedings of the 41st Annual Hawaii International Conference on System Sciences:157. IEEE Computer Society.
- [12] Giles, J. 2005. Internet encyclopaedias go head to head. *Nat.* 438(7070):900–901.
- [13] Ito, J. 2004. Wikipedia attacked by ignorant reporter. Joi Ito's Web. <http://joi.ito.com/weblog/2004/08/29/wikipedia-attac.html>. Accessed 3 February 2009.
- [14] Keen, A. 2007. *The cult of the amateur: how today's Internet is killing our culture*. Doubleday Business, New York.
- [15] Kittur, A. and Kraut, R.E. 2008. Harnessing the wisdom of crowds in Wikipedia: quality through coordination. CSCW 2008: Proceedings of the ACM Conference on Computer-Supported Cooperative Work. New York: ACM Press.
- [16] Loubser, M. and den Besten, M. 2008. Wikipedia admins and templates: the organizational capabilities of a peer production effort. SSRN. <http://ssrn.com/abstract=1116171>. Accessed 1 December 2008.
- [17] Potthast, M., Stein, B. and Gerling, R. 2008. Automatic vandalism detection in Wikipedia. In: Advances in information retrieval, C. Macdonald, I. Ounis, V. Plachouras, I. Ruthven, R.W. White, Eds., Boston, Springer:663–668.
- [18] Ren, Y., Kraut, R. and Kiesler, S. 2007. Applying common identity and bond theory to design of online communities. *Org. Studies*, 28(3):377–408.
- [19] Sanger, L. 2001. Wikipedia is wide open. Why is it growing so fast? Why isn't it full of nonsense?. Kuro5hin. <http://www.kuro5hin.org/story/2001/9/24/43858/2479>. Accessed 10 November 2008.
- [20] Schindler, M. and Vrandečić, D. 2009. Introducing new features to Wikipedia. In: Proceedings of WebSci'09, Athens, Greece.
- [21] Thom-Santelli, J., Cosley, D.R. and Gay, G. 2009. What's mine is mine: territoriality in collaborative authoring. In: Proceedings of the 27th international conference on Human factors in computing systems:1481–1484, Boston, MA, USA. ACM.
- [22] Viégas, F., Wattenberg, M. and Dave, K. 2004. Studying cooperation and conflict between authors with history flow visualizations. In: Proceedings of the SIGCHI Conference on Human factors in computing systems 2004, New York, ACM Press: 575–582.
- [23] Viégas, F., Wattenberg, M., Kriss, F. and van Ham, F. 2007. Talk before you type: coordination in Wikipedia. In: Proceedings of the 40th Annual Hawaii International Conference on System Sciences. IEEE Computer Society.
- [24] Wikipedia:How to write Simple English articles. 2008. In: Simple Wikipedia, the free encyclopedia. http://simple.wikipedia.org/wiki/Wikipedia:How_to_write_Simple_English_articles. Accessed 10 November 2008.
- [25] Wikipedia:Simple English Wikipedia. 2008. In: Simple Wikipedia, the free encyclopedia. http://simple.wikipedia.org/wiki/Wikipedia:Simple_English_Wikipedia. Accessed 10 November 2008.