

Identifying, Understanding and Detecting Recurring, Harmful Behavior Patterns in Collaborative Wikipedia Editing – Doctoral Proposal –

Fabian Flöck
Institute AIFB
Karlsruhe Institute of Technology
Germany
fabian.floeck@kit.edu

Supervised by Dr. Elena Simperl, University of Southampton, UK
and Dr. Achim Rettinger, Karlsruhe Institute of Technology, Germany

ABSTRACT

In this doctoral proposal, we describe an approach to identify recurring, collective behavioral mechanisms in the collaborative interactions of Wikipedia editors that have the potential to undermine the ideals of quality, neutrality and completeness of article content. We outline how we plan to parametrize these patterns in order to understand their emergence and evolution and measure their effective impact on content production in Wikipedia. On top of these results we intend to build end-user tools to increase the transparency of the evolution of articles and equip editors with more elaborated quality monitors. We also sketch out our evaluation plans and report on already accomplished tasks.

Categories and Subject Descriptors

H.5 [Information Interfaces and Presentation]: Group and Organization Interfaces—*Collaborative computing, Computer-supported cooperative work, Web-based interaction*; H.1 [Models and Principles]: User/Machine Systems—*Human factors*

Keywords

Wikipedia, editing behavior, user modeling, collaboration systems, collective intelligence, web science, social dynamics

1. PROBLEM STATEMENT

In sociology, economy and psychology, there have been decades of research on the emergence of harmful social interaction patterns in certain populations, leading to a vast body of scientific work aiming to understand, predict, and prevent the occurrence of such phenomena. These include mass hysteria and panics, stock market bubbles and disease spread, to name only a very few.¹ By contrast, little is known about how such damaging social dynamics form in collaborative online environments that potentially bring

¹A plethora of research work exists on these and many related topics. In the scope of this PhD proposal we cannot introduce them in their entirety. See [19] (panics), [21] (stock market) and [6] (disease spread) as examples.

together orders of magnitude more users than their offline counterparts. These environments thrive and prosper from the contributions of their user base, and their success hinges greatly on the extent to which interaction and collaboration among their users leads to a critical mass of useful, high-quality content.

One of the most prominent examples of online collaborative platforms in the past decade is undoubtedly Wikipedia. It is the direct result of the actions of a community of volunteers who create and edit its articles; members of this community and the systematic behavioral patterns emerging from their interactions form, directly or indirectly, intentionally or unintentionally, the decision about which topics are included in the encyclopedia and how they are represented. In order for such “wisdom-of-the-crowds”-driven judgements to have the desired effects, and to propel the creation of high-quality articles - complete, factually correct and unbiased - it is essential to study and analyze the underlying community and its actions, and to devise methods to identify frequently occurring social phenomena that are likely to indicate a dysfunctional style of collaboration. Ownership behavior, for instance, wherein an individual editor or a group of editors defend an article from being changed by third parties has been identified in the literature, as we will see in Section 1.1. Arguably, this behavior has ambivalent effects: On the one hand, it protects articles from vandalism, on the other hand, it may discourage the introduction of novel information and different points of view by outsiders. There are many other behavioral patterns studied in classical social sciences which can easily emerge in online settings, causing (or facilitating) such harmful impairments. Note that by these behavioral patterns we do not refer to purposeful malicious behavior like vandalism or other actions by individual users, but collaborative interaction patterns of conduct in the (well-intended) editing process.

Understanding and systematically detecting social mechanisms within Wikipedia is a challenging topic; while empirical studies attest the existence of a variety of recurring interaction patterns, the effects of such interaction often remain unclear, and imprecision and hidden assumptions may lead to unwanted conclusions. This is partially due to the fact that assessing the quality of Wikipedia content is highly contextual; there is not, and will never be, a consistent alternative source of information, accepted by all relevant stakeholders, and containing all facts, opinions and viewpoints that should be objectively included in a Wikipedia article. If automatic techniques are available, they are naturally restricted to specific sources, which the content of articles is compared against,

thus identifying only a share of potential inconsistencies or missing information. Such methods typically identify major quality flaws in articles, but are less reliable for cases in which such issues are slightly less obvious, such as a biased point of view. Aside from purely content-driven methods, means to detect the emergence of harmful editor behavior patterns promise a complementary approach; once such behavior occurs, the system generates warnings, and the community at large can react and assess the current status of an article apart from simple fact checking, sentiment analysis and biased wording.

Supplying readers and editors with instruments to make the “social evolution” of an article transparent and explicit can be an effective tool to induce awareness of possible incorrectness and bias and counter uncritical perception. As Wikipedia matures, so need the methods to secure its quality. Giving Wikipedians (and any Wiki users) more elaborated tools to see complex problematic patterns in editor behavior is one step in this direction. Our work aims to provide such methods, which are accurate regarding the way information is processed and interpreted, and representative with respect to the communities of readers and editors they claim to apply to.

The Wikimedia foundation has reached similar conclusions and incorporated them into their strategy to advance the Wikipedia platform. The research planned in this thesis is carried out in the context of the RENDER project², in collaboration with the German chapter of the Wikimedia foundation. Results and tools will be deployed on Wikipedia’s sites and promoted by Wikimedia to the worldwide community.³

1.1 Related work

The work presented in the following exemplifies that Wikipedia articles are increasingly showing traits that hint at the emergence of exactly the harmful behavior patterns we aim to detect. A great majority of the work mentioned here offers descriptive views on single phenomena the authors observed; as a whole, these studies draw a consistent picture of the effects recurring behavior patterns can have on articles, and on the potential of these effects to raise quality and neutrality issues in Wikipedia content.⁴

Some findings [12] suggest that a certain share of legitimate viewpoints⁵ might not be represented in Wikipedia because active editors cover a narrow section of the offline-population’s socio-demographic scope, a problem the community has long identified and labeled the “systemic bias” of Wikipedia.⁶

For those active in the system, the widely observed phenomenon of a small minority of users providing most of the edits and content [25, 14] to an online collaboration system holds true for Wikipedia as well. Priedhorsky et al. [20] show that a vast majority of the content that is actually provided by a minuscule fraction of all editors. This effect is reinforced by the fact that high-frequency editors continue to increase their proportional number of edits [22].

Suh et al. [22] suggest that the declining growth of the English Wikipedia indicates a state of maturity where many articles are close to complete on a factual level. Correlating with this development, for many articles a consolidated article text has emerged since Wikipedia’s inception, which now is relatively fixed insofar

that it is hard for new and occasional editors to change content. As [22] point out, the ratio of reverted edits to the total number of edits has increased with occasional editors experiencing greater resistance compared to high-frequency ones.⁷ Relevant to this scenario [13] show that if an editor tries to remove words from an article, the probability of her edit to be reverted increases significantly with the number of article revisions the removed words had ‘survived’ before (normalized for the number of removed words). In short, words which have been in the article for a longer period of time are harder to remove, as more trust is placed in older content. Likewise, viewpoints added in the early days of Wikipedia could have a much higher probability of being eventually represented in the article, supported by the observations of Kittur et al. [14, 15].

The cause of these phenomena seems to lie in the perceived consensus between the editors, which serves as “social proof” [5] of the correctness of the information. If a majority has (implicitly) assessed a content to be right, it is difficult to change it. An “information cascade” [3] can occur where all participants place their decisions on the preceding editors’ choices instead of their own knowledge. All in all, this valid social heuristic may also hamper the replacement of outdated content or the revision of biased content towards more balanced opinion expressions, at least for newcomers or anonymous users.

Another observation is that certain editors tend to become over-protective about article content against changes from other editors. Although explicitly discouraged by Wikipedia,⁸ strong feelings of ownership towards an article and protective behavior are not uncommon and “[...] run the risk of deterring new community member participation.” [23]. Halfaker et al. [13] infer that “[...] Wikipedia’s review system suffers from a crucial bias: editors appear to inappropriately defend their own contributions.” Articles guarded in such a way naturally run the risk of being biased as new contributions tend to get accepted only if conforming to the owner’s taste. Members or subgroups which do not agree with the majority decision might at first fight for their viewpoint(s) but then eventually just stop editing altogether. Halfaker et al. [13] find indicators for the drop-out of highly reverted editors in the first 16 weeks of membership in Wikipedia. Another effect is the emergence of user ‘camps’ with divergent positions engaging in revert wars on highly controversial articles, which has been investigated in [17]. The case of Conservapedia⁹ reveals that an opposing camp might leave the discussion arena altogether, perhaps resulting in a highly biased article.

The phenomena indicated by the research work presented above make it seem likely that social barriers exist for new viewpoints to be accepted in Wikipedia, even when these may objectively contain useful information. As content matures, and the share of “hard” factual information to be added to articles is diminishing, the value and accuracy of new revisions (for example, the rewrite of an article to include a different point of view on the topic) are far more challenging to evaluate and judge. This means that more complex social control mechanisms and procedures come into play, for instance, in form of a stronger emphasis on reputation and social proof, which may deter some editors from contributing effectively. At the same time, tools and methods to detect the emergence of those possibly negative habits in an article become essential given the scale and complexity of the Wikipedia environment.

²<http://www.render-project.eu>

³See, e.g., the WIKIGINI tool presented in section 2.5, which is already included in the diversity toolkit by Wikimedia.

⁴Note that all of the research presented here was done using data and observations based on the English version of Wikipedia, a practice that we will adopt for the work outlined in the following.

⁵“Legitimate” meaning all content apart from vandalism.

⁶http://en.wikipedia.org/w/index.php?title=Wikipedia:Systemic_bias&oldid=516438857

⁷Indicators such as page protection, deletion, block, and other restrictive policies exhibit a similar trend.

⁸http://en.wikipedia.org/wiki/Wikipedia:Ownership_of_articles

⁹The ultra-conservative “sister” of Wikipedia, founded by ex-Wikipedians (<http://www.conservapedia.com>).

2. APPROACH, STATE-OF-THE-ART AND CONTRIBUTIONS

The planned doctoral thesis aims to (i) give a comprehensive account of the core behavioral mechanisms that fundamentally drive social interaction phenomena leading to biased, incorrect or incomplete content; (ii) systematically identify and categorize these mechanisms; (iii) evaluate what their actual (negative) effects are likely to be on the resulting content under given circumstances; and finally (iv) generate user-friendly support tools and interfaces that inform editors about these phenomena, generating transparency regarding the article evolution and giving an opportunity to take counteractive measures.

The research approach can be divided into five tasks as follows:

Task 1 Systematic collection of patterns

Task 2 Data mining and preparation

Task 3 Collecting a ground truth

Task 4 Pattern and effect analysis

Task 5 Building end-user tools

2.1 Task 1: Systematic collection of patterns

As a first step, it is important to identify those fundamental behavioral mechanisms that have the potential to do notable harm in the collaboration process and to the resulting content, in order to make clear what patterns to look for in the existing Wikipedia edit log data. This was already done to an extent by the author on the one hand through related work showing the damaging effects of imitation behavior in collaborative online systems [7], and on the other hand through a first literature review [10] and initial systematization of recurring suspected harmful behaviors, similar to the overview given in Section 1.1. This review will be expanded to give an account of all relevant empirical research work.

2.1.1 *State-of-the-art*

There is plenty of research work that identifies individual phenomena in Wikipedia, as introduced in Section 1.1. Few of these publications, however, go beyond a descriptive account of the happenings; explanations of the underlying behavior causing these developments or means to measure them are largely missing.

As an exception, e.g., Suh et al. [22] explain the slowing growth and increasing conflict by comparing Wikipedia to a kind of information ecosystem which is overcrowded and cannot “support” its inhabitants anymore (in line with the observations of [4, 11, 17]). Based on their empirical findings, some works claim that tasks requiring low coordination between participants actually profit from many contributors [26], whereas more complex tasks with high coordination overhead, such as building the structure of an article, are optimally accomplished within a comparatively smaller group of users and would be probably unfeasible otherwise [15, 16, 26].

Apart from a few attempts, though, the research on Wikipedia editing dynamics so far fails to provide a comprehensive catalogue of the underlying, recurring mechanisms leading to the phenomena described in Section 1.1, and does not evaluate them in terms of their actual effects on the content quality of Wikipedia.

2.2 Task 2: Data mining and preparation

As a basis to systematically (and even automatically) spot these behavioral patterns in the Wikipedia edit data, meaningful relations between the editors as well as reliable information regarding the authorship of words in the articles have to be extracted from the raw

editing logs provided by Wikipedia. This data is needed for the detection of nearly all relevant dynamics identified in Task 1. The first step was completed in [9], where we devised a sufficiently accurate method to detect reverts between editors. It uses a definition of reverts rooted in the official Wikipedia definition and daily interactions of editors we observed. The evaluation showed a significant increase in accuracy predicting antagonistic relationships between edits compared to previous methods.

Furthermore, we improved the accuracy of the detection of authorship of words in Wikipedia [8] compared to existing methods. This is a task still in progress until completely satisfactory accuracy performance is reached.

2.2.1 *State-of-the-art*

For revert detection, the available research uniformly considers reverts simply as “going back to a previous revision”, leading to straightforward and fast methods, but neglecting the full Wikipedia definition of reverts and the correct interpretation of antagonistic edit behavior [13, 14, 17, 20, 22, 18]. This caveat was removed by our work on reverts.

Few approaches exist for identifying authorship of words. One is *HistoryFlow* by IBM, which operates on a sentence level [24], as well as some minor tools by the Wikipedia community, which show no competitive performance in terms of speed or accuracy. A more elaborated approach is used by *Wikitrust* [1, 2]. It operates on a word level and can as well track reintroduced words and assign the original author. An evaluation showed that a first version of our algorithm outperformed Wikitrust by 10% in terms of accuracy [8]. As of writing this paper, the accuracy of our method lies at circa 90%, thereby surpassing the state of the art by over 50%.

2.3 Task 3: Collecting a ground truth

In order to decide if certain behavior patterns lead to bias or otherwise low quality, it is necessary to determine if an article exhibits these traits in the first place. Using only the mark-up (in the form of tags and ratings) provided by the Wikipedia community will fall short of taking into account the assessment of the shown content by the majority of read-only users that don’t cross the threshold to editorship or users in spe who don’t yet have access to Wikipedia. I.e., relying only on the ratings of the content by the individuals that actually created the content will not be sufficient for an objective analysis. More importantly, most articles in Wikipedia lack *any* metadata indicating their quality, especially crucial in cases of missing tags for low quality which can lead to misplaced trust in the content. Therefore, we will use crowdsourcing techniques to perform content analyses through a diverse set of non-wikipedia-users. This will include (i) testing the ability of workers on crowdsourcing platforms to spot different types of impairments of article content and finding the right setting parameters to optimize the reliability of the results and (ii) extracting the actual ratings from the crowd to be used as a gold standard in the remaining tasks of this doctoral thesis.

2.3.1 *State-of-the-art*

To evaluate if an article is biased or what quality level it has, available research work almost uniformly relies on the ratings given to articles by the Wikipedia community itself, either in the form of specific tags assigned to them by individual editors and wiki projects or, much less frequently, some (semi-)external ratings, e.g. by the Wikipedia 1.0 project or external experts.¹⁰ We are not

¹⁰In the latter case, this usually includes not more than about 20 manually evaluated articles.

aware of any work that tries to collect ratings for Wikipedia articles from external, independent sources and avoiding self-selection problems, especially not from crowdsourcing platforms, making use of the immense diversity of workers they offer.

2.4 Task 4: Pattern and effect analysis

Once collected and identified in Task 1, recurring editing mechanisms suspicious of harming article content will be made detectable by learning which typical structures they exhibit in regard to the social network between users and the network relating users and words in an article, as well as distributions of edits and authorship concentration, revert rates of newcomers and many other variables, all of which will be derived from the data mined in Task 2. From the analysis of such statistical and social network data, measures will be obtained that can indicate the presence and intensity of a given mechanism, e.g., ownership behavior. The methodology used will incorporate traditional statistical methods like regression analysis, clustering and factor analysis as well as social network methods, including exponential random graph models and block modelling.

The next step will be to assess the effect of one or a combination of these patterns in an article on the resulting quality and bias, using the ground truth generated in Task 3. We will look at the correlation between the two and conduct a careful analysis on the possible causalities.

2.4.1 State-of-the-Art

The statistical methods described above have been traditionally applied in social sciences, and social network analysis methods are gaining more and more momentum in the realm of online collaboration, proving their usefulness for gaining deeper insights into this kind of data. We will adopt these to our specific use case.

2.5 Task 5: Building end-user tools

The goal of our work is ultimately to provide users with intuitive tools and visualizations that make possibly negative developments in an article's quasi-hidden social context easy to spot and understand. This holds true for the most recent, ongoing process of content generation, but also for the historic view on the evolution of an article. The latter can yield insights into the quality and neutrality status of the article that are hard or impossible to assess only by reading the article content and comparing it to external sources.

To do so, we will provide tools that are centered around the development over time of those critical indicators (identified in Task 4) that have been singled out as influential for the unfolding of an article, and highlight important or unexpected events in the evolution of these measures, such as spikes or sudden drops (in the simplest case). Furthermore, this information has to be put in relation to other article-related features such as edits and views received, tags and templates added, or activities on the associated talk page. Such tools could be used to augment the outcomes of content-analysis methods, for example algorithms examining fact coverage in an article against external sources such as news or social media.

We have already developed the WIKIGINI tool, which, based on the results of Task 2, determines the concentration of authorship in an article, and its oscillation over time.¹¹ First results of the ongoing analysis we are running with the help of this rather simple indicator are promising in respect to its ability to facilitate the understanding of the evolution of articles, and the identification of critical developments that deserve special attention.

¹¹<http://toolserver.org/~RENDER/toolkit/WIKIGINI/>

2.5.1 State-of-the-art

A variety of tools has been developed for Wikipedia to provide an overview on the editing dynamics for different purposes. One of the earliest visualization tools was *HistoryFlow*, mentioned in Section 2.1. By tracking how long certain words remain unrevised in an article, the *Wikitrust* Project by Adler et al. [1] uses visual markups for "trusted" and "untrusted" passages in any Wikipedia article text in different color shades. Community extensions, like *WikiPraise*¹² expand its original functionality. It does however not give insights as to the diversity of the article content. Further, the definition of "trust" implies quality, when, as a matter of fact, consolidated article content can as well hint at outdated or biased content, as we discussed in section 1.1. *Wikidashboard*¹³ provides an aggregated view of edits in an article performed over time for the most active users. There is, furthermore, a wide variety of tools created by the Wikipedia community.¹⁴ These solutions, however, have neither the purpose nor the ability to uncover neutrality issues or to detect harmful behavioral patterns and generate warnings or at least an easily understandable score for a certain quality or neutrality thread.

3. EVALUATION

Every step in the outlined research methodology will be evaluated to ensure the accuracy and usefulness of the results of each task. This is even more important as the Tasks build up on each other.

Task 1 can not be evaluated in a strict sense, but an extended literature review will be compiled and submitted as a journal article, covering as much empirical evidence as possible for potential damaging behavior patterns in Wikipedia.

Task 2 has been evaluated for the revert detection through a user study with 35 Wikipedia editors which yielded significant statical improvements in accuracy [9]. For authorship detection we used 250 words from 45 randomly selected articles and could report a statistically significant 10% gain in accuracy [8]. Further improvements for authorship detection are planned to be tested with even larger samples.

Task 3: The user studies using crowdsourcing platforms will include settings for cross-evaluating the results between different sample groups, controlling for group effects and self-selection, as well as survey variations ruling out suggestive questions. Also, a pre-study will determine if the workers on crowdsourcing platforms are able to solve these kind of rating tasks and how parameters have to be set in terms of payment, length and types of articles and qualification of workers.

Task 4 will be evaluated by testing the prediction capabilities of the learned patterns for biased and other problematic content through probabilistic modeling.

Task 5 is planned to be tested by doing user studies with Wikipedia editors, readers and non-readers to gauge the usability and usefulness of the tools to heighten transparency, spotting critical developments and help finding the cause for content insufficiencies. Part of the tools will be included and evaluated in the frame of the official diversity toolkit by Wikimedia Deutschland, developed in the research project RENDER.¹⁵

¹²<http://de.wikipedia.org/wiki/Benutzer:NetAction/WikiTrust/WikiPraise>

¹³<http://wikidashboard.appspot.com/>

¹⁴Some examples: Wikirage (<http://www.wikirage.com>) shows trending topics and most editing users recently. WikiXRay (<http://meta.wikimedia.org/wiki/WikiXRay>) generates a range of quantitative descriptive statistics.

¹⁵<http://toolserver.org/~RENDER/toolkit/>

4. CONCLUSIONS

In this doctoral proposal, we made apparent that the maturing of Wikipedia in terms of the content and the collaboration habits of its editor base make more sophisticated methods and tools necessary to keep track of possible harmful social interaction patterns. It was outlined how we plan to systematize these patterns from empirical findings, how we are going to mine and enrich editing log data and how we are going to use this data to parametrize the found mechanisms to eventually build tools that can visualize them and generate warnings for editors and readers. In this way we hope to at least create transparency and awareness for the reader regarding these underlying obstacles for quality and at best provide a lever for editors to remove them. For future work, we will focus mainly on the Tasks 3, 4 and 5, as they are still largely unsolved.

Acknowledgements

The research for this thesis has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 257790, project RENDER.²

5. REFERENCES

- [1] T. Adler and L. Alfaro. A content-driven reputation system for the Wikipedia. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 261–270, 2007.
- [2] T. Adler, K. Chatterjee, L. Alfaro, M. Faella, I. Pye, and V. Raman. Assigning trust to Wikipedia content. In *International Symposium on Wikis*, 2008.
- [3] S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, pages 992–1026, 1992.
- [4] B. Butler, E. Joyce, and J. Pike. Don't look now, but we've created a bureaucracy: the nature and roles of policies and rules in wikipedia. In *Proceedings of the CHI '08*, pages 1101–1110, New York, NY, USA, 2008. ACM.
- [5] R. Cialdini. *Influence science and practice*. Gardners Books Ltd, [S.l.], 2007.
- [6] S. Eubank, H. Guclu, V.S.A. Kumar, M.V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429(6988):180–184, 2004.
- [7] F. Flöck, J. Putzke, S. Steinfelds, K. Fischbach, and D. Schoder. Imitation and quality of tags in social bookmarking systems – collective intelligence leading to folksonomies. In T. Bastiaens et al., editor, *On Collective Intelligence*, volume 76 of *Advances in Intelligent and Soft Computing*, pages 75–91. Springer Berlin / Heidelberg, 2011.
- [8] F. Flöck and A. Rodchenko. Whose article is it anyway? – detecting authorship distribution in wikipedia articles over time with wikigini. In *Online proceedings of the Wikipedia Academy 2012*. Wikimedia, 2012.
- [9] F. Flöck, D. Vrandečić, and E. Simperl. Revisiting reverts: accurate revert detection in wikipedia. In *Proceedings of the 23rd ACM conference on Hypertext and social media*, pages 3–12. ACM, 2012.
- [10] F. Flöck, D. Vrandečić, and E. Simperl. Towards a diversity-minded Wikipedia. In *Proceedings of the ACM 3rd International Conference on Web Science 2011*, 06 2011.
- [11] A. Forte and A. Bruckman. Scaling consensus: Increasing decentralization in wikipedia governance. In *HICSS*, page 157. IEEE Computer Society, 2008.
- [12] R. Glott, P. Schmidt, and R. Ghosh. Wikipedia survey — overview of results. Technical report, UNU- MERIT, United Nations University, Maastricht, Netherlands, March 2010.
- [13] A. Halfaker, A. Kittur, R. Kraut, and J. Riedl. A jury of your peers: quality, experience and ownership in wikipedia. In *Int. Sym. Wikis*. ACM, 2009.
- [14] A. Kittur, E. H. Chi, B. A. Pendleton, B. Suh, and T. Mytkowicz. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. *25th Annual ACM Conference on Human Factors in Computing Systems (CHI'07)*, 2007.
- [15] A. Kittur and R. E. Kraut. Harnessing the wisdom of crowds in wikipedia: quality through coordination. In *CSCW '08: Proceedings of the ACM 2008 conference on Computer supported cooperative work*, pages 37–46, New York, NY, USA, 2008. ACM.
- [16] A. Kittur, B. A. Pendleton, and R. E. Kraut. Herding the cats: the influence of groups in coordinating peer production. In *Int. Sym. Wikis*. ACM, 2009.
- [17] A. Kittur, B. Suh, B. A. Pendleton, and E. H. Chi. He says, she says: conflict and coordination in wikipedia. In *Proceedings of the SIGCHI conference on Human factors in computing systems, CHI '07*, pages 453–462, New York, NY, USA, 2007. ACM.
- [18] S. K. Lam, S. Uduwage, Z. Dong, S. Sen, D. R. Musicant, Loren Terveen, and John Riedl. Wp:clubhouse? an exploration of wikipedia's gender imbalance. In *WikiSym 2011*, Mountain View, CA, 10/2011 2011. ACM, ACM.
- [19] A.R. Mawson. Understanding mass panic and other collective responses to threat and disaster. *Psychiatry: Interpersonal and biological processes*, 68(2):95–113, 2005.
- [20] R. Priedhorsky, J. Chen, S. K. Lam, K. Panciera, L. Terveen, and J. Riedl. Creating, destroying, and restoring value in wikipedia. In *GROUP '07: Proceedings of the 2007 international ACM conference on Supporting group work*, pages 259–268, New York, NY, USA, 2007. ACM.
- [21] R.J. Shiller, S. Fischer, and B.M. Friedman. Stock prices and social dynamics. *Brookings Papers on Economic Activity*, 1984(2):457–510, 1984.
- [22] B. Suh, G. Convertino, E. H. Chi, and P. Pirolli. The singularity is not near: slowing growth of wikipedia. In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration, WikiSym '09*, pages 8:1–8:10, New York, NY, USA, 2009. ACM.
- [23] J. Thom-Santelli, D. Cosley, and G. Gay. What's mine is mine: territoriality in collaborative authoring. In *CHI'09*, pages 1481–1484. ACM, 2009.
- [24] F. B. Viégas, M. Wattenberg, and K. Dave. Studying cooperation and conflict between authors with history flow visualizations. In *Proceedings of the SIGCHI conference on Human factors in computing systems, CHI '04*, pages 575–582, New York, NY, USA, 2004. ACM.
- [25] S. Whittaker, L. Terveen, H. Hill, and L. Cherny. The dynamics of mass interaction. In *Proceedings of the 1998 ACM conference on Computer supported cooperative work (CSCW98)*, 1998.
- [26] D. Wilkinson, D.M. Wilkinson, and B.A. Huberman. Assessing the value of cooperation in wikipedia. *Arxiv preprint cs/0702140*, 2007.