

Towards a diversity-minded Wikipedia

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

Denny Vrandečić
Institute AIFB
Karlsruhe Institute of
Technology
Karlsruhe, Germany
denny.vrandecic@kit.edu

Elena Simperl
Institute AIFB
Karlsruhe Institute of
Technology
Karlsruhe, Germany
elena.simperl@kit.edu

ABSTRACT

Wikipedia is a top-ten Web site providing a free encyclopedia created by an open community of volunteer contributors. As investigated in various studies over the past years, contributors have different backgrounds, mindsets and biases; however, the effects - positive and negative - of this diversity on the quality of the Wikipedia content, and on the sustainability of the overall project are yet only partially understood. In this paper we discuss these effects through an analysis of existing scholarly literature in the area and identify directions for future research and development; we also present an approach for diversity-minded content management within Wikipedia that combines techniques from semantic technologies, data and text mining and quantitative social dynamics analysis to create greater awareness of diversity-related issues within the Wikipedia community, give readers access to indicators and metrics to understand biases and their impact on the quality of Wikipedia articles, and support editors in achieving balanced versions of these articles that leverage the wealth of knowledge and perspectives inherent to large-scale collaboration.

Categories and Subject Descriptors

H.5 [Information Interfaces and Presentation]: Group and Organization Interfaces—*Collaborative computing, Computer-supported cooperative work, Web-based interaction*; H.3 [Information Storage and Retrieval]: Online Information Services—*Web-based services*; H.1 [Information Systems Applications]: Models and Principles—*Human factors*

Keywords

Wikipedia, diversity, community-driven content creation, social dynamics, opinion mining, sentiment analysis

1. INTRODUCTION

Twenty years after its introduction, the Web is *the* platform for the publication, use and exchange of information, on a planetary scale, on virtually every topic, and representing an amazing wealth of opinions, viewpoints, mindsets and backgrounds. The success of the Web can be attributed to several factors, most notably to its prin-

ciplined scalable design, but also to a number of subsequent developments such as user-generated content, smart mobile devices, and cloud computing. These trends are said to be responsible for dramatically lowering the last barriers of entry when it comes to producing and consuming information online, leading to an unprecedented growth and mass collaboration. They empowered hundreds of millions of users all over the globe to create high-quality encyclopedias, contribute to world-class software, publish Terabytes of multimedia content on social networking platforms, and lively take part in defining the agenda of many aspects of our society by publicly expressing and sharing their ideas and resources for the collective good [25]. The downside of this unprecedented success story is the great challenges associated with making sense out of the sheer amounts of information continuously being made available online, while leveraging the diversity inherently unfolding through global-scale collaboration [24]. These challenges are still to be solved at many levels, in particular at the level of the ICT solutions enabling any Web-based undertaking, including the infrastructure to store and access the information, the techniques to manage, analyze and use it, and the paradigms underlying the processes of Web-based information provision and consumption. As the Web is growing to permeate every aspect of our lives, the core technological components of the Web ecosystem have become tremendously influential with respect to the ways information is discovered, filtered, processed, and presented. More specifically, the technologies deployed within widely used platforms such as Wikipedia, Facebook, YouTube, Amazon, Google, and the blogosphere as a whole - be that popularity-based ranking, collaborative filtering, or global truth-driven mechanisms for decision making, to name only a few - determine to which extent and how information is delivered to Internet users, significantly influencing public opinion and the formation and evolution of communities and other means of communication and collaboration.

The Web as it is built today facilitates the discovery and exchange of ideas among people with shared interests, and the creation of globally reaching communities. This paradigm works well for a multitude of application scenarios. Nevertheless, as it is based on consensus-finding mechanisms such as the one in Wikipedia, it is less open to a transparent representation and support of the dialogue between a plurality of views and opinions [21], which is characteristic and beneficial for equally many aspects of our lives, most notably for any kind of socio-political deliberation in accordance to the principles and credos of Western democracy. In fact, one could argue that, despite the pivotal effects of Web-based global collaboration, dialogue and opinion forming are essentially broken on the current Web [24]. Members of a community of interest tend to reinforce each other in their points of view, and only

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.

WebSci '11, June 14-17, 2011, Koblenz, Germany.
Copyright 2011 ACM.

those new members who already think alike, or are willing to accept the decisions the group has previously taken, integrate easily. Consequently, it can be observed that such communities often foster an in-group agreed opinion that may (in time) significantly diverge from the opinion of society at large. As such, the society as a whole becomes increasingly polarized, making it almost impossible to discuss topics on a broader, all-society encompassing scale as inter-group communication and collaboration are hindered. There are numerous examples of these effects; this paper focuses on Wikipedia, as one of the most representative showcases of our time for the power and reach of Web-enabled collective intelligence.

The collaboration paradigms behind the information management technology used within Wikipedia, the procedures installed to govern information provisioning, and the information consumption services that can be offered to end-users based on the available technology, hamper the sustainable growth of Wikipedia. The full range of implications of this state of affairs are still subject of ongoing studies and interdisciplinary research, but some effects can already be observed. To pick just one example, Conservapedia¹ is a project started by former Wikipedia contributors who argue that Wikipedia exhibits a bias towards liberalism and atheism. Relatedly, whereas the English Wikipedia remains a common battlefield for supporters of the Serbian and Croatian points of view on many topics of their common history, the Serbian and Croatian language versions of Wikipedia unveil clear and distinct biases, which are growing due to their relatively separated and self-moderating communities. At another level, Wikipedia as a whole has to deal with huge amounts of information, which are for obvious reasons rich in diversity. To do so, Wikipedia editors have to rely on information management tools that are not designed to reflect, analyze and exploit this diversity, and on a gradually decreasing number of voluntary contributors.

In this paper we discuss the effects of diversity on the quality of Wikipedia content and the evolution of its community of contributors. To do so we undertake an analysis of existing scholarly literature in the area and identify directions for future research and development; we also present an approach for diversity-minded content management within Wikipedia that combines techniques from semantic technologies, data and text mining and quantitative social dynamics analysis to create greater awareness of diversity-related issues within the Wikipedia community, give readers access to indicators and metrics to understand biases and their impact on the quality of Wikipedia articles, and support editors in achieving balanced versions of these articles that leverage the wealth of knowledge and perspectives inherent to large-scale collaboration.

2. HOW DIVERSITY-FRIENDLY IS WIKIPEDIA TODAY

Wikipedia is a top-ten Web site providing a community-built encyclopedia for free. Its success hinges on the support of its volunteer contributors. As illustrated through its governance principles and guidelines, one of the central aims of Wikipedia is to provide a balanced coverage and representation of topics of general interest; according to the *Neutral Point Of View* policy set-up at Wikipedia, the aim is to provide an environment in which an open community of contributors collaboratively creates content "*representing fairly, proportionately, and as far as possible without bias, all signifi-*

¹<http://www.conservapedia.com/>

cant views that have been published by reliable sources."² Despite these noble intentions, systemic bias is introduced by the individual views of the actual contributors, the support each side of a (controversial) discussion receives among Wikipedians, and the social and operational procedures in place for the creation and editing of articles. The importance of these issues for the sustainable future of Wikipedia is expressed most prominently in the latest strategic plan released by the Wikimedia foundation, which lists the improvement of article quality and the expansion of its community of contributors in terms of involvement and cultural and gender diversity as two of its five key priorities by 2015.³

In the scientific literature concerned with Wikipedia, there has been a fair share of research regarding social interaction, editing behavior and collaborative content production. Most of this research primarily or exclusively deals with article quality rather than article diversity in terms of the opinions and viewpoints expressed.⁴ Still, taking a 'diversity-minded' perspective on the findings, one can easily build a storyline that gives an informative account on the representativeness and coverage of Wikipedia articles compared to the multitude of opinions and viewpoints voiced elsewhere on the Web, and on the relationships between diversity and specific socio-technical components underlying Wikipedia as a project, its supporting community, and enabling technology and tools.

In the following we will give an overview of collaborative content production within Wikipedia as it has been analyzed and discussed in recent research literature and various surveys in order to derive indicators and identify possible reasons for the diversity-related problems discussed in Section 1. In particular, we are concerned with questions such as the over- and under-representation of viewpoints, the potential difficulties specific opinions experience to be appropriately reflected in the content of a Wikipedia article, and the relationship between specific characteristics of opinion holders and these difficulties, or lack thereof.

2.1 Representativeness of the community of contributors

Some findings suggest that there might exist a certain share of legitimate viewpoints⁵ that are not represented in the Wikipedia because the editors actually contributing most of its content are not socio-demographically or mindset-wise representative of society in general, of the average Internet users, or even the average reader of Wikipedia. For one, the sparse data that is available on the topic of who is actually contributing to Wikipedia suggests that active editors cover a narrow section of the offline-population's socio-demographic scope, a problem known from polling methodology as 'coverage bias' [18]. The UNU-Merit Wikipedia Survey [9], an online survey conducted with over 300.000 Wikipedians at the end of 2008 revealed that, among other results, of all users interviewed, less than 13% of contributors are female, and a mere 9, 2% are regular and 24, 5% are occasional contributors; that some countries such as Germany are vastly over-represented in terms of number

²http://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view

³<http://blog.wikimedia.org/2011/02/25/wikimedia-presents-its-five-year-strategic-plan/>

⁴Note that (if not marked otherwise) all of the research presented here was done using data and observations based on the English version of Wikipedia and the results might not be transferable to other language editions because of different technical features, cultural background, language, etc.

⁵'Legitimate' meaning all content apart from vandalism.

of editors in relation to their population size; that half of the respondents are between 18 and 30 years old and that 48% of those who contribute have tertiary (undergraduate) or higher education. In the light of these statistics, it is safe to assume that if only such a narrow socio-demographic selection of Internet users - or of the worlds population as a whole - passes the threshold to authorship, the viewpoints inherently introduced to the articles via their contributing will be equally representative of the general public or even the common Wikipedia readership. This imbalanced state of affairs is exacerbated by the fact that every individual editor is biased along various dimensions, deliberately because of an agenda or an opinion or unconsciously because she cannot possibly know about every existing viewpoint on a topic. The openness of an online collaboration system such as Wikipedia is a pre-requisite for the diversity of its contributing users, but self-selection can still lead to some form of homogeneity inside the system [17].

Further, for those who actually contribute the widely observed phenomenon of a small minority of users [20, 31] providing most of the edits [27, 29]] and content [12] to an online collaboration system holds true for Wikipedia as well. Priedhorsky et al. [22] analyzed the relationship between content that is actually read and its provenance, showing that on average 44% of an article came from the 0,1% top-frequent editors, and 85% of it from the 10% top-frequent ones, a ratio which even exceeds a power law distribution in terms of elitism of users. This effect is reinforced by the fact that high-frequency editors continue to increase their number of edits, as reported by [23] in an analysis from 2009. [12] see a positive correlation between the experience of the user measured in number of previous edits and the content added per edit. Complementarily the long tail of articles, in other words the topic diversity in Wikipedia, is growing [16], but only according to Wilkinson and Huberman [32] “[...]a small number of articles, corresponding to topics of high relevance or visibility, accrete a disproportionately large number of edits, while the vast majority of articles experience far less activity.” The low-activity articles hence run even a greater risk of not representing a very large proportion of viewpoints as they lack a broad range of contributors. These phenomena narrow down the diversity of viewpoints that could be theoretically acquired via collaborative editing.

A possible explanation as to why this is (necessarily) happening in an environment such as Wikipedia is the need for coordination to achieve a feasible level of group performance in the content production process [11, 13, 14, 30, 32]. [13] argue based on their empirical findings that tasks requiring low coordination between participants profit from many contributors (see also [32]), whereas more complex, high-coordination tasks, such as building the basic structure of an article are optimally done within a comparatively smaller group of users and would be probably unfeasible otherwise.

2.2 Content consolidation and saturation

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 [23] point out, the resistance measured as the ratio of reverted edits to the total number of edits has increased from 2.9% in 2005 to 6% in 2008, with occasional editors experiencing greater resistance compared to high-frequency ones. Indicators such as page protection, deletion, block, and other restrictive policies exhibit a similar trend. This trend towards content consolidation is very important when it comes to answering the question if and how new points of view can be adequately represented, and if and which others are

already manifested in the current version of an article on historical and other grounds. Hence, we take a closer look at indicators and reasons for the consolidation of content, and for the emergence of increasingly higher barriers to newcomers to add (legitimate) content that may lead to diversification.

Although Wikipedia has shown exponential growth in the past [5, 15, 27, 32], since 2007 this trend has been declining in the number of edits and editors per month. As can be seen in Table 1, the absolute number of new articles has been decreasing during the last two years dramatically (by more than 10% per year), a phenomenon most notably visible for the English and German versions of the encyclopedia, which are considered to be 'mature' in terms of overall article count.⁶ The number of edits, however, is either growing (across all Wikipedia's) or stagnating at a high level (for the German and English Wikipedia's), as listed in Table 2.⁷

[23] discuss this decline in the English Wikipedia and explain it by comparing Wikipedia to a kind of information ecosystem which has reached a state of maturity where many articles are close to complete on a factual level. Accordingly they trace back the decline to “(a) the slowing growth of the editor population due to limited opportunities in making novel contributions; and (b) increased patterns of conflict and dominance due to the consequences of the increasingly limited opportunities.” This observation is in line with other works that identified a shift from the number of edits to coordination, policy setting, and governance [4, 8, 15]. It also seems like Wikipedia is 'running out of easy topics', in other words, topics a larger share of casual users can make useful contributions to without being an expert or spending too much effort. In addition, as Wikipedia grows larger, it becomes increasingly difficult to set-up new articles [19].

2.3 Collaborative content production: decision-making procedures and social factors

2.3.1 Consensus: good heuristic for quality, barrier for diversity

Even if topic 'completeness' has not been accomplished yet, in many cases a social consensus has been reached for an article or subsection of an article, meaning that the content has been (re-)viewed by many users and eventually reached a stable state, where the editors so far have brought in their ideas and which has not been changed for a longer time. [10] show, for example, that if an editor removes words in her edit, 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). Adler et al. [1, 2] base their WikiTrust metrics on author reputation derived from the persistence of her revisions and demonstrate that reputation is a good predictor for her later edits to be less likely removed. In both cases, the interpretation of these findings is that word and edit persistence are a measurement of quality according to Wikipedia rules, for instance, a predictor for 'featured articles' [1, 2]. Still, in a nutshell the results of the studies cited in fact make the case for basically one thing: that words which have been in the article for a longer period of time are harder to remove. The reasons for this phenomenon seem to be because of the perceived consensus between the editors - as in: “If those words were read by so many people and were not changed, they

⁶ Figures are based on http://en.wikipedia.org/wiki/Wikipedia:Modelling_Wikipedia_extended_growth.

⁷ Figures are based on <http://stats.wikimedia.org>.

Table 1: Growth rate of the English Wikipedia. Article count on January 1st of each year

Year	Article count	Annual increase	% annual increase	Avg daily increase
2002	19.700	19.700	-	54
2003	96.500	76.800	390%	210
2004	188.800	92.300	96%	253
2005	438.500	249.700	132%	682
2006	895.000	456.500	104%	1.251
2007	1.560.000	665.000	74%	1.822
2008	2.153.000	593.000	38%	1.625
2009	2.679.000	526.000	24%	1.437
2010	3.143.000	464.000	17%	1.271
2011	3.518.000	375.000	12%	1.027

Table 2: Number of articles and their monthly number of edits

Year	Articles (all)	Edits (all)	Articles (English)	Edits (English)	Articles (German)	Edits (German)
2002	21.000	16.000	17.000	14.000	1.100	180
2003	149.000	114.000	104.000	75.000	12.000	16.000
2004	440.000	414.000	196.000	175.000	48.000	61.000
2005	1.400.000	1.600.000	446.000	618.000	194.000	318.000
2006	3.300.000	6.000.000	895.000	2.700.000	360.000	727.000
2007	6.200.000	10.500.000	1.500.000	4.500.000	563.000	960.000
2008	9.400.000	10.900.000	2.100.000	4.300.000	732.000	827.000
2009	12.200.000	11.600.000	2.700.000	4.300.000	892.000	878.000
2010	14.900.000	11.800.000	3.200.000	4.000.000	1.000.000	845.000
2011	17.900.000	12.800.000	3.600.000	4.100.000	1.200.000	870.000

certainly have to be right”, which serves as ‘social proof’ [6] of the correctness of the information. If a majority has (implicitly) assessed a content to be right, it is difficult to change it. All in all, this is a valid and necessary social heuristic when trying to reach and maintain high-quality articles, as it successfully blocks vandalism and unqualified contributions. In the same time, however, it also hampers the replacement of outdated content or the revision of biased content towards more balanced opinion expressions, at least for newcomers or anonymous users (see also Section 2.3.2). In addition, viewpoints added in the early days of Wikipedia could have a much higher probability of being eventually represented in the article because of information cascades, as has been argued in a different context in [7], and supported by the observation of Kit-tur et al. [12] that “Just as the first pioneers built infrastructure which diminished future migration costs, the early elite users of Wikipedia built up enough content, procedures, and guidelines to make Wikipedia into a useful tool that promoted and rewarded participation by new users”, confirmed in later studies [13]. While efforts are being made to alleviate this situation, as illustrated, for instance, in the strategic goals of the Wikimedia foundation by 2015, the question which still remains to be answered is how much of the article content and the socio-technical infrastructure set-up in the early days of Wikipedia will persist, and how the previous versions have influenced what readers see today. With these findings in mind the assessment of Halfaker et al. [10] that in an ‘ideal system’ (i.e., a high-quality Wikipedia) removing established words should be positively correlated to the probability of being reverted is not to be supported from a diversity point of view.⁸

⁸This partial mismatch between the requirements for quality and diversity is as well reflected by the fact that the community might vote a high-consensus article ‘featured’ (to be read as ‘high-quality

2.3.2 Boldness and the usefulness of conflict

Boldness, conceived as challenging old content structures by bringing in new content, is important to achieve greater diversity, and actively encouraged in Wikipedia.⁹ Very similar to this notion [15] point out that that specific kinds of conflict can be beneficial in peer collaboration systems, while [28] make a congruent point for the case of conflicts which are not related to personal attributes of the contributors, but the content. On a related note, Wilkinson and Huberman [32] recommend increasing the number of users with diverse viewpoints to a high-conflict article instead of having “the same few people arguing back and forth.” as a means to avoid deadlock situations and achieve progress. Nevertheless, according to Bryant et al. [3] editors usually get bolder as they gain more editing experience, meaning that many occasional contributors might not reach a level of boldness to enter in confrontations to get their viewpoints accepted.

2.3.3 Opinion camps and editor drop-out

Related to the potential drawbacks of a ‘fixed consensus’ is the observation that 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.[10] find indicators for the drop-out of highly reverted editors in the first 16 weeks of their contributing to Wikipedia. This could be, on the one hand, due to the fact that their work does not meet the quality criteria required by Wikipedia. On the other hand, it seems likely that

by Wikipedia rules’), while this can still mean that it excludes points of view, as the authors participating in the vote might not represent or be aware of divergent viewpoints.

⁹http://en.wikipedia.org/wiki/Wikipedia:Be_bold

many authors leave because their point of view is not taken into account to the desired degree. This is the case for the previously mentioned Conservapedia, a clone of Wikipedia which represents views on many topics that are typically referred to as 'conservative'. A second effect of this phenomenon is the emergence of user 'camps' with divergent positions engaging in revert wars on highly controversial articles, which has been investigated in [15]. The example of Conservapedia reveals that such opposing camps might not in every case re-enter the discussion arena to eventually find a consensus with the rest of the community, but that one of the parties win the fight entirely, resulting in a highly biased article. This 'tyranny of the majority' is a known characteristics of deliberative democracy, where the majority puts their interest above that of the minorities, and of various online replicas thereof [17], which in Wikipedia's case would mean a violation of its *No democracy* policy.¹⁰

2.3.4 *Wikipedia as a bureaucracy and low motivation of new editors to contribute*

The success of Wikipedia has been accompanied by a massive extension of its governance apparatus, leading, among other things, to additional overhead in activities related to the creation, maintenance and enforcement of the associated rules and policies [4, 8, 15]. While this can be seen as a sign of maturity and progress, the complexity of explicit as well as implicit social norms can be perceived as negative and demotivating by newcomers striving to comprehend the rationales for their contributions not being acknowledged in the form of actual changes to the content of Wikipedia articles of their interest [23]. Unexperienced editors may find it difficult to find their way in the bureaucracy [19]. Increasing the barrier of entry for new contributors by the introduction of a stricter, extended governance system is reinforcing the biases that are likely to plague a considerable share of the Wikipedia content (in particular, beyond the well-developed versions in mainstream languages) which can be traced back to the characteristics of the group of actual contributors, and the constraints imposed by official regulations and content production processes and practices.

2.3.5 *Territoriality and ownership*

Although explicitly discouraged by Wikipedia,¹¹ strong feelings of ownership for an article and protective behavior are not uncommon [26]. This might have to do with the hours of work many authors put into an article, the self-perceived level of expertise they possess on a given topic, and other reasons that lead to a personal attachment to an article. Although Thom-Santelli et al. [26] highlight the positive effects of a territorial watch for deterring vandalism, they "[...] also observe that these defensive behaviors may run the risk of deterring new community member participation." Halfaker et al. [10] show that the number of editors whose words are reverted during an edit (normalized for the number of deleted words) is a very strong and stable predictor of the probability that the deleting edit will be reverted itself. They find this 'stepping on toes' effect to be in place independently of any other feature of the editor and hence 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.

¹⁰<http://en.wikipedia.org/wiki/Wikipedia:Democracy>

¹¹http://en.wikipedia.org/wiki/Wikipedia:Ownership_of_articles

2.3.6 *Personal characteristics of the authors*

Halfaker et al. [10] point out that "If persistent properties of editors, such as knowledge, skill or personality, are related to the quality of an editor's work, then we should see that the probability of a revert is a property of an editor[...]" and report that rarely reverted editors will continue to be rarely reverted in future revisions. They relate the probability of the survival of an edit to features of the editor such as the quality of her work, but also to discussion skills, endurance, or the sheer amount of time she's willing to spend on defending her viewpoint. If the most active and most prominent (by viewed words) editors share such characteristics, it is not unlikely that some of those properties correlate with a specific bias as well. Insightful research on this topic is not available yet to our knowledge, but would be very promising as an indicator for biases and imbalance. Another threat to diversity could be the self-selection of authors. Although one of the major advantages of open online collaboration systems is the self-allocation of an editor to his topics of interest and self-perceived level of expertise this might lead to "insufficient diversity in points of view" [15]. For example, the editors with the highest interest could end up editing the most for a topic (with a certain shared viewpoint), as it is very likely for e.g., topics with an important meaning to religious groups while still being of broad public interest.¹²

2.3.7 *WikiProjects and other groups*

Related to the notions of self-selection and territoriality discussed above is the effect of WikiProjects, where editors group up to care for articles inside a specific topical domain.¹³ Users tend to shift their editing to project-related articles upon joining a project and start engaging more in discussion on the respective talk pages [14]. Similarly, they seem to internalize the goals (and probably views) of the project, which then might also be reflected in the way they edit. The internalization of project viewpoints on certain topics might as well hamper the diversity of viewpoints expressed by project members on that topic through their edits.

2.4 Summary

The phenomena discussed in this section make it seem likely that high barriers exist for new viewpoints to be accepted in Wikipedia, even if they objectively contain useful information. As content matures, and the share of non-contested 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 control mechanisms and procedures come into play, for instance, in form of more explicit rules and a stronger emphasis on reputation and social proof, which may deter some editors from contributing effectively. Also, as defending one's edits gets increasingly difficult, personal features of the editors such as durability might become more important to defend a revision. At the same time, editors who have put a considerable amount of effort in their articles and have a strong personal motivation to contribute tend to be protective about what they have crafted over a long period of time. The same protectiveness could be attributed to WikiProjects, which follow a strict agenda.

Researchers addressed these issue by designing tools that aim to enhance the understanding of collaborative editing and decision-making processes, and their implications from a content and a community-

¹²See, for instance, the article on evolution at <http://en.wikipedia.org/wiki/Talk:Evolution>.

¹³<http://en.wikipedia.org/wiki/WikiProject>

oriented perspective within the Wikipedia environment. Adler et al. [1] propose a system that computes quantitative values of trust for the text of Wikipedia articles by visualizing the trust of individual terms computed on the basis of the reputation of the associated editor. A second example along the same lines is WikiDashboard, a social dynamics analysis tool which improves social transparency by making the hidden social context of pages and articles of Wikipedia explicit.

3. LEVERAGING DIVERSITY IN WIKIPEDIA

One way to tackle the issues constraining viewpoint diversity is to revise the existing Wikipedia processes. Activities such as resolving conflicts, the organization of the content, the checking of inconsistencies and biases (both within Wikipedia and with respect to external sources), and the integration over different languages (for checking fact coverage and biases) are very demanding in terms of the amounts of human labor, transparency and coordination they require. Elaborated procedures that cover some of these aspects, including edit conflict resolution, arbitration committees, and banning policies, and a growingly complex hierarchy of readers, contributors, editors, administrators, bureaucrats, ombudspersons, trustees, and so on, are in place, but their operation, given a declining number of active Wikipedians and the complexity of the tasks, is not sustainable. Also, the outcomes of these costly processes are not always positive; the meritocratic approach of Wikipedia often finds champions for specific opinions, but seldom for a generally balanced, diversity-minded depiction of a topic.

Together with the German chapter of the Wikimedia Foundation, the European research project RENDER¹⁴ will work on building a truly diversified Wikipedia. Wikipedia editors need support in discovering useful content and the diversity of viewpoints within a topic to encourage large-scale participation and sustainable growth. Using the massive amount of metadata available within Wikipedia (which directly scales with the number of edits shown in Table 2), as well as a series of structured and semi-structured external information sources, we will provide representations, techniques and tools to discover, understand, and use the following types of information: the multitude of opinions, sentiments and viewpoints, the points of dissent, content that would otherwise disappear from view, the quality of articles, and controversies surrounding specific topics. Information sources that are useful in this context include, but are not limited to, the complete edit history of each article, change comments, user contribution logs, the implicit social networks in the user contribution logs - as in, who works with or against whom on which articles - the content of articles, including comments and previous versions, access logs and various external data sets such as Eurostat, data.gov, Twitter, Linked Open Data, Freebase, Wolfram Alpha, and archives of scientific publications. In order to provide these representations, techniques and tools, we will tackle a number of challenges:

Understand and predict socio-technical mechanisms leading to biases

For relevant phenomena potentially leading to imbalances in coverage, representation and accuracy of information (some of them presented in this paper) we will develop models to show how they function, which effects on diversity they actually have and which patterns they display that can be used to detect and predict them. As a first step we will look into the existence and effects of territoriality and ownership behavior in Wikipedia and to other possible effects of heightened

vigilance e.g., triggered by vandalism attacks. Specifically, we want to know if the ability of new and occasional editors to add their points of view to the article is impaired by those effects. The analysis incorporates several features of authors, articles and edits taken from the complete revision history of the articles and their discussion pages.

Identify and extract diversely expressed information We will provide mechanisms to identify and extract opinions, viewpoints and sentiments based on the available Wikipedia metadata, going significantly beyond shallow text mining and information extraction. These mechanisms will use concepts represented in the article text, their relation to each other, temporally coincident comments on the article talk pages, and answers to these, as well as discussions on the responsible contributors' talk pages.

Represent and process diversely expressed information We will design methods that utilize opinions and viewpoints to summarize, understand, and visualize the flow of discussions on a specific topic. As a highly expressive formalization of discussions cannot be achieved in a feasible way - due to the limitations of formal knowledge representation languages paired with the computational complexity associated with inferring over such rich formalizations - our methods will leverage semi-structured data such as fragments of articles and associated change information, as well as lightweight representations and reasoning that make key aspects of diversity explicit.

Building upon the results of these three lines of research we aim to support a number of diversity-empowered services for

Displaying warnings when detecting patterns of bias We will collect and analyze sources about a topic, and compare the fact coverage in the external sources to the facts covered in the Wikipedia article about the topic. Wikipedia exists in more than 250 language editions. We want to compare the different language editions of an article to discover if certain language editions expose biases. When such a situation occurs, the system will display an appropriate warning to allow readers to understand that an article is biased.

Supporting readers to extend articles Whenever our system discovers that an article is lacking a certain point of view, we can offer editors links to sources or extracted summaries of the missing facts. This will support authors to achieve an unbiased presentation of the topic in an article. It also raises awareness among editors for different points of view.

Detecting change in the topic of an article We will monitor external news streams in order to detect certain types of events. This will enable us to discover that an article or a part of an article may be outdated. Within RENDER, we will also analyze the news streams for their biases and will thus be able to exploit this information to provide a balanced presentation of changes to the Wikipedia editors.

The result of the confidence and diversity analysis should be made accessible and explicit to Wikipedia readers. Currently, manual tagging of articles exist, but the tags are often out-of-date and incompletely applied to the article set. Automatic tagging mechanisms will provide the reader with more confidence about the level

¹⁴<http://www.render-project.eu>

of bias of the article. Besides tagging complete articles, users should also be able to mark a single statement and query the system about that statement. Often, Wikipedia articles contain in general good knowledge, but are sprinkled with small inconsistencies or simple acts of vandalism. Whereas blunt lies are often discovered and corrected quickly, subtle errors may escape the attention of most readers and linger in the article for a long time. Our tools will learn to categorize and understand edits to Wikipedia, and record this information as metadata to each article. All metadata that will be collected during the project will be made available to the community, so that further tools can be developed in order to assess the quality of an article or article revision.

4. CONCLUSIONS AND FUTURE WORK

We gave an overview of problems identified in relation to biases and diversity in Wikipedia and their presumed causes. For tackling the diversity issues and some other problems of Wikipedia, we proposed three use cases which will be developed and implemented together with the German Wikimedia Chapter using a combination of semantic representations and reasoning, text and data mining and statistical analysis of socio-technical aspects of content production in Wikipedia. As a first step of our research we will first try to tackle the problem of understanding and quantifying the effect of territorial behavior regarding content leading to biases, and produce prototypes for fact coverage and bias detection on article text.

5. ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 257790.

6. 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. L. Bryant, A. Forte, and A. Bruckman. Becoming Wikipedian: Transformation of Participation in a Collaborative Online Encyclopedia. In *GROUP '05: Proceedings of the 2005 international ACM SIGGROUP conference on supporting group work*, pages 1–10, 2005.
- [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 *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1101–1110, 2008.
- [5] A. Capocci, V. D. P. Servedio, F. Colaiori, L. S. Buriol, D. Donato, S. Leonardi, and G. Caldarelli. Preferential attachment in the growth of social networks: The Internet encyclopedia Wikipedia. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, 74(3), 2006.
- [6] R. Cialdini. *Influence science and practice*. Gardners Books Ltd, 2007.
- [7] F. Floeck, J. Putzke, S. Steinfels, K. Fischbach, and D. Schoder. Imitation and Quality of Tags in Social Bookmarking Systems – Collective Intelligence Leading to Folksonomies. In T. Bastiaens, U. Baumöl, and B. Krämer, editors, *On Collective Intelligence*, pages 75–91. 2011.
- [8] A. Forte and A. Bruckman. Scaling consensus: Increasing decentralization in wikipedia governance. In *HICSS*, page 157, 2008.
- [9] R. Glott, P. Schmidt, and R. Ghosh. Wikipedia survey — overview of results. Technical report, UNU- MERIT, United Nations University, Maastricht, Netherlands, Mar. 2010.
- [10] A. Halfaker, A. Kittur, R. Kraut, and J. Riedl. A jury of your peers: quality, experience and ownership in wikipedia. In D. Riehle and A. Bruckman, editors, *Int. Sym. Wikis*, 2009.
- [11] G. Kane. It's a network, not an encyclopedia: A social network perspective on wikipedia collaboration. In *Academy of Management Annual Meeting*, 2009.
- [12] 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 2007)*, 2007.
- [13] 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.
- [14] A. Kittur, B. A. Pendleton, and R. E. Kraut. Herding the cats: the influence of groups in coordinating peer production. In *International Symposium on Wikis*, 2009.
- [15] 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.
- [16] S. K. Lam and J. Riedl. Is Wikipedia growing a longer tail? In S. D. Teasley, E. C. Havn, W. Prinz, and W. G. Lutters, editors, *GROUP*, pages 105–114, 2009.
- [17] D. Lemire, S. Downes, and S. Paquet. Diversity in open social networks. Technical report, University of Quebec, Montreal, CA, Oct. 2008.
- [18] P. S. Levy and S. Lemeshow. *Sampling of populations*. Wiley series in probability and statistics. Wiley, New York, NY [u.a.], 3. ed edition, 1999.
- [19] A. i. Lih. The battle for wikipedia's soul. *The Economist*, Mar. 2008.
- [20] K. Marks. Power laws and blogs, 2003.
- [21] K. O'Hara and D. Stevens. The devil's long tail: Religious moderation and extremism on the web. *IEEE Intelligent Systems*, 24:37–43, 2009.
- [22] R. Priedhorsky, J. Chen, S. T. K. Lam, K. Panciera, L. Terveen, and J. Riedl. In *GROUP '07: Proceedings of the 2007 international ACM conference on Supporting group work*, pages 259–268.
- [23] 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, 2009.
- [24] C. Sunstein. *Republic.com 2.0*. Princeton University Press, 2007.
- [25] D. Tapscott and A. D. Williams. *Wikinomics: How Mass Collaboration Changes Everything*. Portfolio, 2006.
- [26] J. Thom-Santelli, D. Cosley, and G. Gay. What's mine is mine: territoriality in collaborative authoring. In D. R. O. Jr, R. B. Arthur, K. Hinckley, M. R. Morris, S. E. Hudson, and

- S. Greenberg, editors, *CHI*, pages 1481–1484, 2009.
- [27] J. Voss. Measuring wikipedia. In *Proceedings of the International Conference of the International Society for Scientometrics and Informetrics (ISSI)*, Stockholm, 2005.
- [28] B. Vuong, E. Lim, A. Sun, M. Le, and H. W. Lauw. On ranking controversies in wikipedia: models and evaluation. In *WSDM '08: Proceedings of the international conference on Web search and web data mining*, pages 171–182, New York, NY, USA, 2008. ACM.
- [29] J. Wales. Jimmy wales talks wikipedia, Dec. 2005.
- [30] H. T. Welsler, D. Cosley, G. Kossinets, A. Lin, F. Dokshin, G. Gay, and M. A. Smith. Finding social roles in wikipedia. In *iConference*, pages 122–129. ACM, 2011.
- [31] 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.
- [32] D. Wilkinson, D. Wilkinson, and B. Huberman. Assessing the value of cooperation in wikipedia. *Arxiv preprint cs/0702140*, 2007.