

User Contribution and Trust in Wikipedia

Sara Javanmardi, Yasser Ganjisaffar, Cristina Lopes and Pierre Baldi

School of Informatics & Computer Sciences

University of California, Irvine

{sjavanma, yganjisa, lopes, pfbaldi}@ics.uci.edu

Abstract— Wikipedia, one of the top ten most visited websites, is commonly viewed as the largest online reference for encyclopedic knowledge. Because of its open editing model –allowing anyone to enter and edit content– Wikipedia’s overall quality has often been questioned as a source of reliable information. Lack of study of the open editing model of Wikipedia and its effectiveness has resulted in a new generation of wikis that restrict contributions to registered users only, using their real names. In this paper, we present an empirical study of user contributions to Wikipedia. We statistically analyze contributions by both anonymous and registered users. The results show that submissions of anonymous and registered users in Wikipedia suggest a power law behavior. About 80% of the revisions are submitted by less than 7% of the users, most of whom are registered users. To further refine the analyzes, we use the Wiki Trust Model (WTM), a user reputation model developed in our previous work to assign a reputation value to each user. As expected, the results show that registered users contribute higher quality content and therefore are assigned higher reputation values. However, a significant number of anonymous users also contribute high-quality content. We provide further evidence that regardless of a user’s attribution, registered or anonymous, high reputation users are the dominant contributors that actively edit Wikipedia articles in order to remove vandalism or poor quality content.

I. INTRODUCTION

Wikipedia is the largest encyclopedia on the Web which is based on crowdsourcing, the process of outsourcing a task to a large group of people, in the form of an open call [1], [2]. The crowdsourcing approach stands in sharp contrast to more traditional models of content creation and publication, which tend to limit content creation to a relatively small group of approved editors in order to exercise strong quality control. Because of its open editing model –allowing anyone to enter and edit content– Wikipedia’s overall quality has often been questioned [3], [4]. Lack of study on the open editing model of Wikipedia has resulted in a new generation of wiki-based encyclopedias like Citizendium¹ and Scholarpedia² which provide restricted platforms allowing only experts to contribute using their real names, and ensuring that each contribution undergoes a peer review process before being published. The emergence of these wikis is based on the assumption that the open editing model of Wikipedia cannot be trusted for creation of high-quality content [4].

In this paper, we challenge this assumption by presenting an empirical study of user contributions to Wikipedia. First,

we describe a statistical analysis aiming to compare the contributions of anonymous and registered users to Wikipedia. We show that registered users contribute to Wikipedia more actively compared to anonymous users. They do most of the corrections while most of the vandalistic content is associated with anonymous users. To further refine the analysis, we use the Wiki Trust Model (WTM) presented in [5] to estimate the reputation of users on a large sample of Wikipedia. The results show that there is a positive correlation between user registration and quality of contributed content. However, a significant number of anonymous users are among the most productive contributors of high-quality content. We provide further evidence that regardless of a user’s attribution, registered or anonymous, high reputation users are the dominant contributors; they actively edit Wikipedia articles in order to remove vandalism or poor quality content.

The remainder of this paper is organized as follows: Section II describes related work. In Section III, we statistically compare the contributions of anonymous and registered users to Wikipedia. In Section IV, we estimate the reputation of anonymous and registered users and analyze their behavior. Finally, in Section V, we draw the main conclusions and suggestions for future research.

II. BACKGROUND AND RELATED WORK

According to the open editing model of Wikipedia, users can contribute anonymously or with untested credentials. Because of this, the quality of Wikipedia’s content has often been questioned since its inception. While it is difficult to measure quality of users’ contributions to Wikipedia, a few studies have tried to make this assessment using manual comparative reviews [6], [7], or more indirect statistical analyzes [8] as well as historical analyzes of contribution patterns [9], [10], [11].

In [6], [7], the authors tried to assess the quality of Wikipedia articles manually by comparing them to their parallel articles in other reputable sources and found no major differences. Likewise, Nature Magazine’s own comparative analysis of forty-two science articles in both Wikipedia and the Encyclopedia Britannica revealed a surprisingly small difference in quality. While useful, it is clear that manual analysis cannot scale up and that more statistical, data mining, approaches are needed to assess the quality of Wikipedia and its contributors on a large scale. In this vein, a limited statistical analysis of Wikipedia [8] showed that the featured,

¹<http://en.citizendium.org/>

²<http://www.scholarpedia.org/>

presumably high-quality articles are distinguishable by their large number of edits and contributors.

Several studies have tried to assess the quality of Wikipedia’s content and the reputation of its contributors by analyzing historical contribution patterns[5], [10], [12], [13]. For instance, Zeng et al. [13] discussed a method to compute the trustworthiness of Wikipedia articles from a dynamic Bayesian network. They hypothesized that the trustworthiness of the revised content of an article depends on the trustworthiness of the previous revision, the authors of the previous revision, and the amount of changes made to the text. To estimate content quality, they categorized users into four groups –administrators, anonymous users, registered users, and blocked users– and assigned a single *static* reputation to each group. Adler and Alfaro [10] assigned a *dynamic* reputation to users based on their actions. In their approach, reputation assessment is based on the survival of the contributed text, and the survival of the contributed edits. Similarly, in our previous work [5], we presented WTM, a trust model that assigns dynamic reputation to users based on their past actions in the system. WTM is an event based model. Whenever a new revision is submitted to the wiki, its text is compared with the text of the previous revisions of the corresponding article in order to extract the properties of newly inserted and deleted content. Contributors whose content remains in the wiki pages gain reputation. Conversely, those whose contributed content is deleted are penalized. The reputation penalty is based on several factors including the size of the deleted content, the reputation of the user who is deleting the content, and the number of revisions that have occurred between the insertion and deletion of the content.

III. STATISTICAL COMPARISON OF WIKIPEDIA USERS

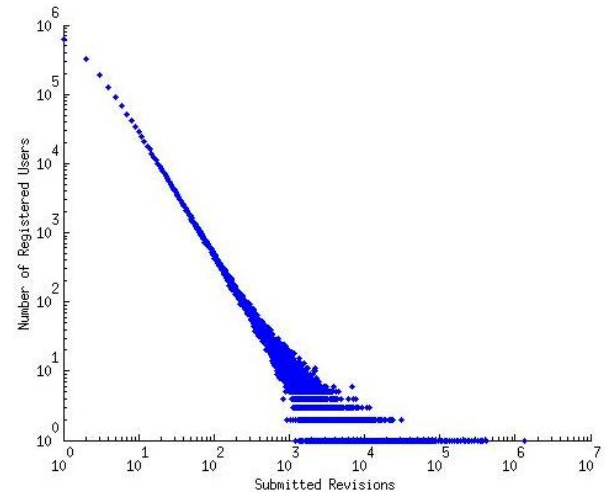
In this section, we present a statistical comparison of the behavior of anonymous and registered users in Wikipedia. According to Wikipedia, registered users are identified by their usernames, while anonymous users are tracked by their IP addresses. As a result of the Internet’s very nature, there is no one-to-one correspondence between people and accounts or IP addresses. We, as well as others [13], follow the same nomenclature as Wikipedia: a “user” in this study refers to a registered account or an IP address, and it does not refer to a real-world individual.

For the purpose of this study, we used the latest dumps of the English Wikipedia released in October 2008³, which contain the history of Wikipedia since its inception in January 2001 (see Table I). The dump also includes information about pages that are not relevant to this study, namely talk pages, image description pages, user profiles, templates, help pages, portals and pages for Wikipedia’s administration. As such, we have excluded them from our study.

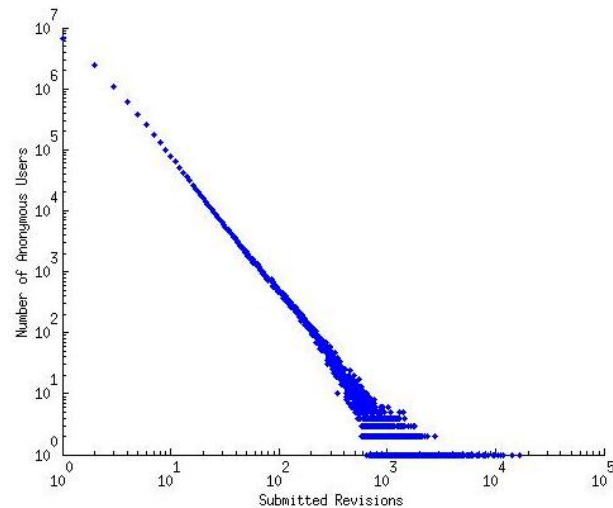
Figure 1 plots the number of registered/anonymous users versus the number of submitted revisions on a log-log scale. The submitted revisions from both anonymous and registered

TABLE I
PROPERTIES OF THE DATA SET

Time span	2,844 days
Number of users	11,250,440
Number of articles	2,448,058
Number of revisions	143,559,522



(a) Registered users



(b) Anonymous users

Fig. 1. Distribution of submitted revisions in Wikipedia.

users suggest a power law behavior. To verify this, we used the method introduced in [14] and the companion software. Our results show that distribution of submitted revisions by registered users (Figure 1(a)) fits a power law distribution with goodness-of-fit value of 0.0045. The goodness-of-fit is calculated by comparing observed data to the hypothesized power law distribution based on the Kolmogorov-Smirnov statistic. This value is 0.0029 for anonymous users (Figure 1(b)). Values less than 0.01 indicate a strong power law behavior.

Table II shows the distribution of submitted revisions over population size. To have more accurate results, we have ex-

³<http://download.wikipedia.org/enwiki/20081008/>

TABLE II
DISTRIBUTION OF SUBMITTED REVISIONS OVER POPULATION SIZE

Revisions	Population Percentage	Revisions by Anonymous vs. Registered Users
20%	0.007%	Anon: 0.03%, Reg: 19.97%
40%	0.05%	Anon: 0.45%, Reg: 39.55%
60%	0.39%	Anon: 4.56%, Reg: 55.44%
80%	6.87%	Anon: 16.06%, Reg: 63.94%
100%	100%	Anon: 33.11%, Reg: 66.89%

cluded revisions submitted automatically by Wikipedia Bots⁴. 80% of the revisions are submitted by less than 7% of the users, while 63.94% of them are contributions of registered users. According to this observation, a small group of users in Wikipedia have submitted most of the revisions.

To compare the behavior of anonymous and registered users in more detail, Table III summarizes the statistics pertaining to users' behavior:

- *Population size* shows the percentage of different groups of users in Wikipedia. Registered users constitute 17% of the population and about 0.01% of these users are Admins and Bots, users with specific privileges.
- *Contribution ratio* shows the percentage of submitted revisions by each group of users. Registered users have contributed 67% of total revisions. Interestingly, Admins and Bots who comprise a small portion of Wikipedia users have submitted 10% and 5.5% of the revisions, respectively.
- *Reverted ratio* shows the percentage of submitted revisions that are reverted. A revert is an action to undo all changes made to an article, restoring it to what it was at a specific time in the past. These numbers do not include self-reverts, i.e. when a user reverts his own revisions to a previous version that he had contributed. According to the Wikipedia's revert policy⁵, reverting is used primarily to fight vandalism or similar activities such as spamming. Revisions submitted by anonymous users constitute 66% of the reverted revisions.
- *Correction ratio*. Reverts are corrective actions which are used by users for speedy removing of vandalism, poor-quality data and points of view from articles. When a user reverts revision r_i to revision r_{i-k} , a new revision, r_{i+1} , is created which contains exactly the same content as r_{i-k} and the user who has done the revert is considered as the author of this new revision. Correction ratio measures the percentage of revisions similar to r_{i+1} in the total revisions submitted by a user.

Although anonymous users comprise 83% of users in Wikipedia, their submission ratio is only 33%. In addition, most of the reverted revisions are associated with anonymous users while registered users do most of the reverts (corrections).

⁴<http://en.wikipedia.org/wiki/Wikipedia:Bots>

⁵http://en.wikipedia.org/wiki/Wikipedia:Revert_only_when_necessary

IV. BEHAVIORAL COMPARISON OF WIKIPEDIA USERS

In the previous section, we presented a statistical analysis of the behavior of anonymous and registered users in Wikipedia. However, not all of the users in the registered or anonymous populations behave similarly. Hence, more fine-grained analysis of the users in each of these two populations can result in a better understanding of Wikipedia. To further the analysis, here we need to dynamically assess the behavior and quality of contributions of each user at each time, by estimating their *reputation* at each time based on their edit patterns over a previous period of time. Intuitively, high-reputation users are those who contribute high quality content; these high quality contributions tend to survive in the articles for a longer time [12], [15]. More precisely, here we use the WTM [5] implementation to dynamically assign a reputation value, between 0 and 1, to each user.

Assume that user i has inserted $N_i(t)$ tokens in the system before time t and $n_i(t)$ of these tokens are not deleted yet. At time t , he inserts $c_i(t)$ new tokens where $g_i(t)$ of them remain in the wiki page while the rest are deleted by other users. Reputation of user i is updated based on the following formula:

$$R_i^+(t) = \max \left(0, \frac{n_i(t) + g_i(t) - \sum_{d=1}^{p_i(t)} R_{j(t_d)} e^{-\alpha(\Delta r)}}{N_i(t) + c_i(t)} \right) \quad (1)$$

where $R_{j(t_d)}$ is the reputation of the deleter at the time of deletion, $p_i(t)$ is the number of deleted tokens, and Δr is the number of revisions submitted between insertion and deletion of the tokens.

Although reputation updates in WTM can be done in a constant time after a revision is submitted, given the number of submitted revisions in Wikipedia since its inception, estimating reputation of all Wikipedia users at all times is computationally expensive. Thus, for the present study, we restricted ourselves to a representative sample.

To generate a representative sample, we modeled Wikipedia as a bipartite-weighted graph. Users and articles are mapped to the graph nodes; contributions are mapped to the weighted edges, where the weight represents the total number of revisions submitted by the user to the article. We used several graph sampling algorithms including Random Walk, Random Edge, Random Jump and Forest Fire [16] in order to find

TABLE III
STATISTICAL COMPARISON OF DIFFERENT USER GROUPS

	Population Size	Contribution Ratio	Reverted Ratio	Correction Ratio
Anonymous Users	83%	33%	66%	2.9%
Registered Users	17%	67%	34%	8.2%
Admins	0.014%	10%	4.9%	18.6%
Bots	0.004%	5.5%	0.7%	6.9%
Other	16.98%	51.5%	28.4%	7%

a suitable sample. In our experiments, the Random Edge algorithm resulted in the sample most similar to the original Wikipedia data set. We measured similarity of samples to the entire Wikipedia data set using several factors including the proportional size of anonymous and registered populations and percentage of contributions by each population. The final sample used corresponds to about 20% of the original Wikipedia data set and is representative of the corresponding bipartite graph [16].

A. Results

Figure 2 shows the distribution of reputation values for anonymous and registered users. It shows that, on average, registered users benefit from higher reputation values. For reputation values less than 0.5, the percentage of anonymous users is greater than that of registered users. Conversely, the percentage of registered users is greater when reputation values above 0.5 are considered. It also shows that a significant percentage of anonymous and registered users lies on the two extremes of the reputation spectrum; there is a significant number of users in both groups, but high-reputation users outnumber low-reputation users.

1) *Contributions Analysis*: To analyze contributions by high and low reputation users to Wikipedia, Figure 3 shows the distribution of user contributions as a function of user reputation. We classify anonymous and registered users into 10 groups based on the reputation values and calculate the contribution ratio for each group. The contribution ratio is defined as the number of submitted revisions by users in the group to the total number of revisions. As the diagram shows, regardless of a user’s attributions, there is a positive correlation between user reputation and number of contributions: users with higher reputation contribute more actively to Wikipedia than users with lower reputation. It also shows that, on average, high-reputation registered users have a higher contribution ratio than high reputation anonymous users.

2) *Revert Analysis*: WTM assigns low-reputation values to users who contribute short-lived content in the articles. Short-lived content is expected to be of poor-quality or associated with vandalism. According to Wikipedia, reverting is used primarily to fight vandalism, or anything similar to the effects of vandalism⁶. We extracted reverts to analyze behavior of each user based on the percentage of his revisions that were

reverted by other users (*i.e.* reverted ratio), and also the percentage of revert actions that he performed (*i.e.* correction ratio).

To extract reverts, the text of each revision in an article is compared to the text of the previous revisions in order to check if they are the same. We compare the MD5 signatures of the texts in order to check content equality more efficiently. The MediaWiki API⁷ is used to download the text of past revisions of each article. Downloading text of past revisions of all Wikipedia articles and extracting reverts from them was done with a cluster of 25 clients over a period 5 days. The results show that 9% of the submitted revisions have been reverted.

Figure 4 shows the distribution of correction and reverted ratios as functions of user reputation. As user reputation increases, correction ratio rises and revert ratio falls.

These results suggest that high-reputation users are the major source of corrections in Wikipedia (reverts). It is important to note that reverts are not include in the WTM inputs, hence the positive correlation between user reputation values and correction ratio, and the negative correlation between user reputation values and reverted ratio provide independent evidence of the accuracy of the WTM model at predicting the reputation of a user consistently with the user’s behavior.

To see how quickly reverts happen in Wikipedia, we also extracted the number of revisions between submission of a revision and its reversion. Since different articles have different visit rates and reverts might happen faster in articles with higher visit rates, we measure time by the number of revisions that occur in it, rather than its continuous value. Our study shows that 75% of the reversions occur immediately by being applied to the immediate previous revision. In 94% of the times, reversions are made to one of the 4 most recent revisions. Since reverts are primarily used to fight vandalism, this result shows that vandalism elimination happens within the next few revisions. Some other studies have also reported quick elimination of vandalism in Wikipedia [17], [18], [19]. Magnus in [19] showed that about one third to one half of the systematically inserted fictitious claims in Wikipedia were corrected within 48 hours.

⁶http://en.wikipedia.org/wiki/Revert_vandalism

⁷<http://en.wikipedia.org/w/api.php>

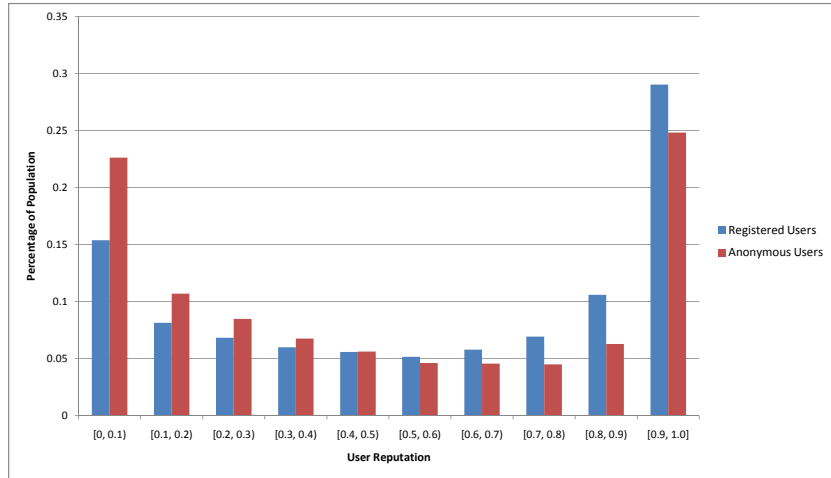


Fig. 2. Distribution of anonymous & registered users over user reputation

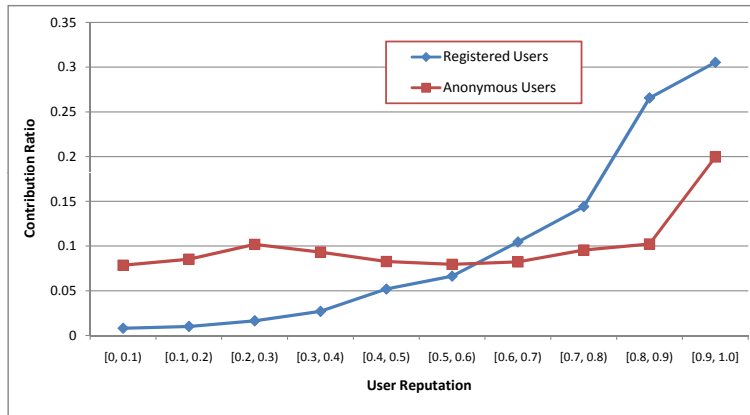


Fig. 3. Distribution of Contributions over user reputation. Contribution Percentage shows the percentage of revisions submitted by each group of users.

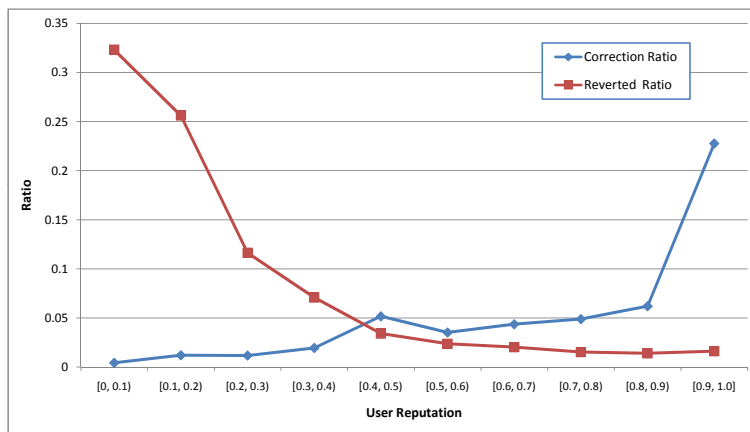


Fig. 4. Distribution of reverts over user reputation

V. CONCLUSION & FUTURE WORK

This study provides an analysis of the open editing model of Wikipedia. We compared the behavior of anonymous and registered users from different aspects. The results show that the majority of the revisions are submitted by a small percentage of users, and that most of them are registered users. Furthermore, the results show that there is a positive correlation between user registration and the quality of contributed content. The distribution of user reputation in Wikipedia also shows that, regardless of attribution, vandals and inexperienced users are still quite a minority compared to high reputation users. Thus, it appears that the open editing model of Wikipedia relies on a large number of well-intentioned users who actively contribute in such a way as to neutralize any vandalism or poor quality content that may be inserted by the smaller number of problematic users. In further analysis of the open editing model of Wikipedia, we aim to compare anonymous and registered users according to the amount of contributed content that has survived in Wiki entries. Although the results so far show that only 6% of the content in Wikipedia belongs to anonymous users, distribution of survived content ratio for anonymous users indicates that a significant fraction of these users contribute high quality content.

ACKNOWLEDGEMENT

This work has been partially supported by NSF grant OCI-074806.

REFERENCES

- [1] A. D. W. Don Tapscott, *Wikinomics: How Mass Collaboration Changes Everything*. Portfolio Hardcover, Dec 2008, pp. 70–77.
- [2] M. Waldrop, “Big data: Wikinomics,” *Nature News*, vol. 455, pp. 22–25, Sep 2008.
- [3] J. Giles, “Wikipedia rival calls in the experts,” *Nature*, vol. 443, pp. 493–493, Oct 2006.
- [4] Editorial, “Wouldn’t you like to know? a wealth of information is available online, in useful encyclopaedic form. but how much of it is to be trusted?” *Nature Physics*, vol. 4, pp. 505–505, Jul 2008.
- [5] S. Javanmardi and C. Lopes, “Modeling trust in collaborative information systems,” in *CollaborateCom '07: Proceedings of the 3rd International Conference on Collaborative computing: Networking, Applications and Worksharing*. New York, NY, USA: IEEE, November 2007, pp. 299–302.
- [6] J. Giles, “Internet encyclopaedias go head to head,” *Nature*, vol. 438, pp. 900–901, Dec 2005.
- [7] T. Chesney, “An empirical examination of wikipedia’s credibility,” *First Monday*, vol. 11, no. 11, 2006. [Online]. Available: http://www.firstmonday.org/issues/issue11_11/chesney/index.html
- [8] D. M. Wilkinson and B. A. Huberman, “Cooperation and quality in wikipedia,” in *WikiSym '07: Proceedings of the 2007 international symposium on Wikis*. New York, NY, USA: ACM, 2007, pp. 157–164.
- [9] M. Kramer, A. Gregorowicz, and B. Iyer, “Wiki trust metrics based on phrasal analysis,” in *Proceedings of the 2008 International Symposium on Wikis (WikiSym) Porto, Portugal*, Sep 2008.
- [10] T. B. Adler and L. de Alfaro, “A content-driven reputation system for the wikipedia,” in *WWW '07: Proceedings of the 16th international conference on World Wide Web*. New York, NY, USA: ACM Press, 2007, pp. 261–270.
- [11] E. H. Chi, B. Suh, and A. Kittur, “Providing social transparency through visualizations in wikipedia,” in *Proceedings of the Social Data Analysis Workshop at CHI'08*, Florence, Italy, Apr 2008.
- [12] M. Hu, E.-P. Lim, A. Sun, H. W. Lauw, and B.-Q. Vuong, “Measuring article quality in wikipedia: models and evaluation,” in *CIKM '07: Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*. New York, NY, USA: ACM, 2007, pp. 243–252.
- [13] H. Zeng, M. Alhossaini, R. Fikes, and D. McGuinness, “Mining revision history to assess trustworthiness of article fragments.” IEEE Computer Society, 2006.
- [14] A. Clauset, C. R. Shalizi, and M. E. J. Newman, “Power-law distributions in empirical data,” *ArXiv e-prints*, vol. 706, Jun 2007.
- [15] B. T. Adler, D. L. Alfaro, I. Pye, and V. Raman, “Measuring author contributions to the wikipedia,” Tech. Rep., Sep 2008.
- [16] J. Leskovec and C. Faloutsos, “Sampling from large graphs,” in *KDD '06: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. New York, NY, USA: ACM, 2006, pp. 631–636.
- [17] F. B. Viégas, M. Wattenberg, and K. Dave, “Studying cooperation and conflict between authors with history flow visualizations,” in *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*. New York, NY, USA: ACM Press, 2004, pp. 575–582.
- [18] A. Kittur, B. Suh, B. A. Pendleton, and E. H. Chi, “He says, she says: conflict and coordination in wikipedia,” in *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*. New York, NY, USA: ACM Press, 2007, pp. 453–462.
- [19] P. Magnus, “Early response to false claims in wikipedia,” *First Monday*, vol. 9, no. 13, 2008.