

Wikipedian Self-Governance in Action: Motivating the Policy Lens

Ivan Beschastnikh

Computer Science and Engineering
University of Washington
ivan@cs.washington.edu

Travis Kriplean

Computer Science and Engineering
University of Washington
travis@cs.washington.edu

David W. McDonald

Information School
University of Washington
dwmc@u.washington.edu

Abstract

While previous studies have used the Wikipedia dataset to provide an understanding of its growth, there have been few attempts to quantitatively analyze the establishment and evolution of the rich social practices that support this editing community. One such social practice is the enactment and creation of Wikipedian policies. We focus on the enactment of policies in discussions on the talk pages that accompany each article. These policy citations are a valuable micro-to-macro connection between everyday action, communal norms and the governance structure of Wikipedia. We find that policies are widely used by registered users and administrators, that their use is converging and stabilizing in and across these groups, and that their use illustrates the growing importance of certain classes of work, in particular source attribution. We also find that participation in Wikipedia's governance structure is inclusionary in practice.

Introduction

In online communities, technology is the medium through which members act – and for many communities, this data is available for researchers to data mine. Despite the availability of such data, it is difficult for quantitative research to move from basic measurements of frequency counts and network representations to confident assertions about the establishment and evolution of social structure. And while qualitative studies can provide rich descriptive information, it is difficult to generalize results. With this paper, we open a dialogue about how quantitative methods might be used to bracket the rich social practices.

We examine Wikipedia, a prominent example of a community seeking to collectively produce and maintain a valuable artifact – a neutral account of all human knowledge. At the heart of this effort is the idealized decision-making mechanism of consensus. As Wikipedia grows in popularity and notoriety, achieving consensus is a challenge given the wide range of backgrounds, intent and attitude of participants. A rich set of policies and guidelines mediate the difficult work of seeking consensus (Kriplean *et al.* 2007; Viégas *et al.* 2007; Forte & Bruckman 2008). This *policy environment* articulates strategies of action, principles of encyclopedic content, and proper user behavior. Each policy

is described and summarized on its own page just like any other Wikipedia article. For example, the *Assume Good Faith* guideline states that users should regard other users' actions as well intentioned, while the *Verifiability* policy states that all facts need to be attributed to a verifiable source. The policy environment encodes and explains norms, but it is not imposed from up high; rather, policies are created and managed by the Wikipedia community at large. This is facilitated by two technical mechanisms:

First, contributors can easily hyperlink to policies during discussions concerning article content. We call these hyperlinks *policy citations*. By relying on policies, contributors can more easily interpret complex situations and legitimate their actions (Kriplean *et al.* 2007). Policy citations serve to help users educate new users, work through and resolve content disputes, and to deal with trolls (Viégas *et al.* 2007).

Second, policy pages are subject to the same open-editing and consensus-seeking processes as the rest of Wikipedia, giving contributors a participatory role in which experiential knowledge and best practices can be shared (Viégas, Wattenberg, & McKeon 2007). The policy environment therefore is reflective of social practice, not simply prescriptive of social behavior (Forte & Bruckman 2008).

We focus on one concrete manifestation of the enactment and enforcement of the policy environment: policy citations. Specifically, *we provide a quantitative account of the practice of citing policies on the talk pages that accompany each article*. We take the stance that we can learn more about the normative structure of the community by how its members enact policy, rather than examining the development of the policies themselves. The policy lens gives us leverage in identifying the challenges that online communities encounter at different scales of participation, strategies for addressing these difficulties, and empirical data for evaluating the efficacy of these strategies. We demonstrate that policy citations track community concerns and hence provide an interesting normative lens. Our results also corroborate and extend a number of qualitative findings about Wikipedia.

We begin by surveying the complex Wikipedia policy environment. Next we position our research amidst a growing body of prior work on Wikipedia. We then describe our methodology for extracting policy use on discussion pages and present our investigations into policy use. We conclude with a discussion of the implications of our findings.

Anatomy of Wikipedia

In this section we describe relevant aspects of Wikipedia, outlining user categories, the article namespaces that comprise Wikipedia, and the policy environment.

Users. Wikipedia provides different levels of technical privileges to classes of users. We distinguish three user classes: *anonymous* users (logged by their IP address), *registered* users (logged by their unique name handle), and *administrators*. Administrators are granted special privileges such as page protection and user blocking. Registered users become administrators after being nominated for adminship by those who recognize significant contributions and adherence to Wikipedia principles. A panel of administrators then reviews each nomination, examining the candidate's history of edits and reading testimonials from the wider community.

Namespaces. Wikipedia pages are divided into namespaces. Namespaces most relevant to our analysis are the *main* namespace containing Wikipedia articles, the *talk* namespace containing talk pages, and the *wikipedia* namespace where policies and other community pages exist. We call pages in the main namespace *articles* and pages in the talk namespace *talk pages*. Other namespaces include *user* (personal pages for registered users) and *user talk* (messages for the user). See Pentzold & Seidenglanz for a more thorough overview. Whenever one views a page, they are looking at a particular *revision* of the page. A *revert* is a type of revision that swaps the current revisions with some prior revision.

Policy environment. The policy environment, established in February 2002 (one year into the project), delimits proper encyclopedic content, defines acceptable behavior, outlines writing style, establishes legitimate reference sources, and describes formal processes for resolving disputes and sanctioning users.¹ Like the rest of Wikipedia, anyone can edit the policy pages. The policy environment is therefore dynamic and evolves as new policies are created, merged, elaborated upon, and clarified. Although the Wikimedia board sometimes impose policy, the policy environment is most accurately conceived of as being reflective of practice (Forte & Bruckman 2008). As of the writing of this paper, the policy environment comprises 38 *official policies* and 189 *guidelines*. Official policies are the most formal and prominent. Guidelines are less official, although violation of a guideline can carry consequences. We generically use the term *policy* when referring to any official policy or guideline.

Policies are often referenced during discussion on the talk pages. Such references help to socialize and discipline new and deviant participants by reinforcing standards of article content and user conduct (Viégas *et al.* 2007; Kriplean *et al.* 2007). These references are often made by users in attempts to claim legitimate control of an article in order to move the consensus process forward (or to hijack it) (Kriplean *et al.* 2007). An author can reference policy in two ways: (1) a *policy citation* is made by creating a hyperlink to the relevant policy page in the Wikipedia namespace and (2) a *keyword reference* references the policy in plain text. See Figure 2 for

¹The policy environment plays a similar role in Wikipedia as FAQs do in Usenet (Kollock & Smith 1996).

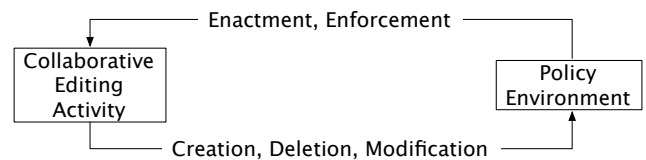


Figure 1: The wiki software is used to create a tight feedback loop between the policy environment and collaborative practice.

U1: Then cite some of those *reliable, notable sources*. That's all I'm asking.

U2: I don't feel its necessary. If you feel its necessary, you are more than welcome to add them.

U1: **"The obligation to provide a reputable source lies with the editors wishing to include the material, not on those seeking to remove it."**

U2: There is nothing more *reputable* than the actual video of [NAME] herself...If you feel more is needed, the burden is on you to add. Please do not delete well-referenced information.

U1: I see we disagree. I assert you are ignoring **WP:V** and **WP:RS** just as you ignore **WP:3RR**. I hope we can resolve these differences but I believe the article is...being led further an further away from the values we hold in Wikipedia (primarily **WP:OR**, **WP:NPOV**, and **WP:V**) and your edits are significantly contributing to the problem.

Figure 2: An anonymized excerpt from a dispute about whether a video clip is a valid reference source. A number of policy citations (bolded) and keyword references (italicized) are made. U1 quotes a sentence from the Verifiability policy and hyperlinks to it. U1 later uses shorthand link to policy (Verifiability [WP:V], Reliable Sources [WP:RS], Neutral Point of View [WP:NPOV], No Original Research [WP:OR], and Three-Revert Rule [WP:3RR]).

an example of a content dispute where participants employ policy citations and keyword references.

Affordances of the wiki software (simple hyperlinking and alternate namespaces) have thus been appropriated to structure collaboration by creating a tight feedback loop between the policy environment and community practice (Figure 1). On one hand, if a policy does not exist or is not well suited to address an issue, members may alter the policy environment itself; on the other, a policy citation brings the policy environment to bear on a specific issue. Enactment of a policy makes that policy visible to other participants, who may also start to enact the policy. If Wikipedians do not actively employ a policy, it ceases to be a structuring feature of the policy environment.

Related Work

We are starting to understand Wikipedia as a rich social space comprised of many intersecting communities (Kriplean *et al.* 2007; Forte & Bruckman 2008). Different forms of work, such as welcoming new users (Viégas, Wattenberg, & McKeon 2007), discussion and consensus seeking on talk pages (Viégas *et al.* 2007; Kriplean *et*

al. 2007), vandal fighting (Viégas, Wattenberg, & Dave 2004), and administrative tasks, are becoming increasingly important. Elaborate graduated dispute resolution processes mediate conflicts between contributors (Forte & Bruckman 2008). Ritual surrounds the granting of administrative privilege. Ideological tensions about the purpose and function of Wikipedia abound (*e.g.* “Inclusionists” and “deletionists” clash over whether Wikipedia’s goal is to be the sum of human knowledge or a carefully delimited encyclopedia).

Technical mechanisms enable different social structuring: policy pages and hyperlinking mechanisms allow users to quickly invoke the voice of the community, “watchlists” allow users to keep track of articles they are interested in (Bryant, Forte, & Bruckman 2005), “recent changes” lists and reverting allow contributors to fight vandalism and spam (Buriol *et al.* 2006), administrative capabilities such as blocking users and removing history have furthered social control, and templates allow detailed categorization of the corpus, such as tagging articles with the necessary work to be done (Cosley *et al.* 2007) or summarizing the most relevant information for classes of articles (Wu & Weld 2007). In this section, we summarize quantitative studies about the Wikipedia and Usenet communities and discuss results of relevant qualitative studies. We end by describing recent work detailing governance in Wikipedia and how our work informs this line of research.

Quantitative. Quantitative studies have established a number of facts about Wikipedia and its community. We know that nearly all measurable quantities yield a skewed distribution (Voss 2005), vandalism (especially mass deletions) are quickly reverted (Viégas, Wattenberg, & Dave 2004), edits to the main article space are proportionally decreasing with respect to the other namespaces (Viégas *et al.* 2007; Kittur *et al.* 2007b), and that featured pages (high-quality articles) are correlated with higher levels of activity by distinct authors (Wilkinson & Huberman 2007), while the number of distinct authors is inversely proportional to the probability that a page is classified as controversial (Kittur *et al.* 2007b). We also know that the top contributors are making proportionally less edits (Kittur *et al.* 2007a; Ortega & Barahona 2007), although they still contribute a large proportion of the valuable text (Priedhorsky *et al.* 2007). Almeida *et al.* (2007) demonstrated that the technical mechanisms that allow Wikipedians to monitor recent changes can have a significant effect on the distribution of edits among articles. Most relevantly, Buriol *et al.* (2006) discovered that policy can strongly affect behavior, finding that reverts were cut in half in the months following the institution of the `three-revert` rule.

Usenet has also received much attention from quantitative researchers. Studies of the social structure of Usenet focused on modeling user and group behavior by creating social networks through metrics such as replies, message length, cross-posting, interactivity, and posting frequencies (Whittaker *et al.* 1998; Fisher, Smith, & Welser 2006; Turner *et al.* 2005). Researchers demonstrated how sequences of actions constitute meaningful roles (Turner *et al.* 2005) and examined how to retain active members (Arguello

et al. 2006). We see Wikipedia as providing the opportunity for quantitative researchers to dig deeper. Datasets such as policy citations open the possibility for enriching models of social structure, the temporal development of communities, and how microactions constitute salient social behaviors.

Qualitative. Using exploratory visualization, Viegas *et al.* (2004) identified a few patterns of activity including vandalism, negotiation, and tensions in authorship, concluding that collaboration in Wikipedia is facilitated by technical and social mechanisms for communal introspection. Although exploratory visualization is somewhat anecdotal, this line of research has uncovered a number of interesting trends (see also (Wattenberg, Viégas, & Hollenbach 2007)). Bryant *et al.* (2005) studied how behaviors and attitudes change as involvement deepens. They found that novice users frame Wikipedia as a collection of articles with random contributors, but as they grow more experienced, they begin to understand Wikipedia as a community with various roles, and subcultures. Viegas *et al.* (2007) studied the functions of talk pages, finding that they are mainly used for coordinating action on the article page. They also found that references to policies and guidelines were present on almost 8% of all talk posts. In a previous study, we analyzed how policies are employed by Wikipedians as they work towards consensus (Kriplean *et al.* 2007). We found that consensus is often moved forward only through a variety of power plays that contributors make in order to claim legitimacy of their actions. Policy citations were used as a sampling strategy to identify potentially interesting social activity.

Governance. Wikipedia’s policy environment is part of a governance structure that has become more elaborate over its history. Recent work has begun to understand this structure in terms of Yochai Benkler’s peer-based commons production model and Ostram’s work on collective self-governance (Viégas, Wattenberg, & McKeon 2007; Forte & Bruckman 2008). This governance structure includes (1) graduated dispute resolution mechanisms (Forte & Bruckman 2008), where violations of policy can be punished through blocking and banning, (2) an elected board of Wikipedians that carries out judicial-like functions such as the interpretation of policy and formal arbitration (Forte & Bruckman 2008), and (3) stringent formal processes for becoming an administrator and elevating a page to “featured” status (Viégas, Wattenberg, & McKeon 2007).

In a study of the featured article process (which dictates how the best Wikipedia articles are selected), Viegas *et al.* (2007) identified that greater emphasis is placed on citing legitimate sources; Forte *et al.* (2008) also find a shifting emphasis towards verifiability of information. We find confirming evidence for this trend in our investigations by noting that policies relating to attribution are more frequently cited than policies in other categories. Forte *et al.* (2008) also find evidence that policy editing is slowing, and the process for adding to the policy environment is becoming more formalized. Later we show that the patterns of policy citations are likewise stabilizing and converging. They also identify that enforcement of policies is becoming more decentralized over time: early on, policy enforcement was per-

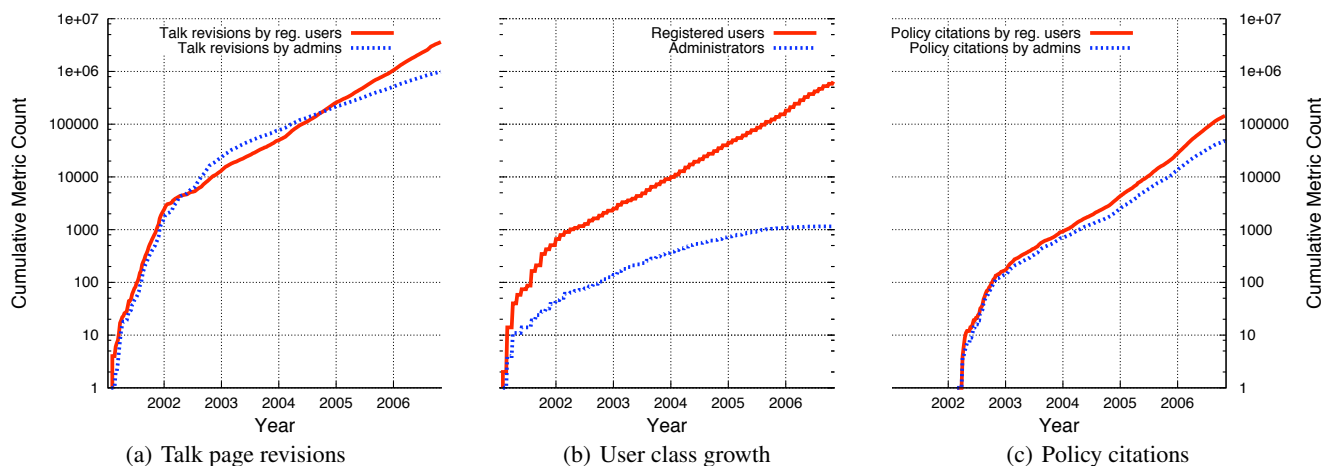


Figure 3: High-level time series trends.

formed by the co-founder, later this responsibility shifted to an elected board of Wikipedians; recently, the task of punishing deviance has been taken up by administrators. Our investigations support this trend.

Our approach differs significantly in emphasis from these interview-based studies that focus on the formal aspects of Wikipedian governance. Policy citations are the everyday manifestation of the governance structure, the means by which individual contributors invoke the norms of the community. An invocation is not a formal process, although it may signal or threaten that a formal action may need to be taken. *We see policy citations as the micro-to-macro link between social norms and the self-governance structure.*

Methodology

The Wikimedia foundation periodically publishes a dataset of all Wikipedia content, including the revision history of pages from all the namespaces. Metadata for each revision is included, containing the username of the editor and the timestamp of the revision. The data we present was derived from the November 4th, 2006 English Wikipedia dataset.

Because Wikipedians consider the talk pages to be their primary communication forum (Bryant, Forte, & Bruckman 2005) we focus on this namespace for our analysis. We mined all policy citations in the talk namespace, specifically all hyperlinks that lead to the main page of a policy, including redirects. For example, WP:NPOV redirects to Wikipedia:Neutral_point_of_view. We also extracted a small set of keywords that unambiguously reference policy. These included clear text statements of “be bold” and “NPOV.”²

Investigations Into Policy Use

To begin exploring the relationship between the policy environment and collaboration in Wikipedia, we present high-

²Edits performed by bots (software agents) are ignored. Archived pages are ignored, except in the cases where the history from the original page was moved to the archive. In this case, the archive and the original page are consolidated.

level time series metrics relevant to the enactment of policy in Figure 3.³ A number of important trends are evident:

- The ratio of policy citations to talk edits is increasing, indicating that the enactment of policy is becoming increasingly common.
- The population of administrators is growing linearly, while the population of registered users is growing exponentially. Despite this trend, administrators are almost maintaining the same growth rate in policy citations as the population of registered users.
- Registered users have overtaken administrators in aggregate policy citations.⁴ This extends the results of Forte *et al.* (2008) by suggesting that enforcement has diffused into the larger body of registered users.

These trends guide our investigation into the establishment and evolution of Wikipedia’s social structure, as reflected by the enactment of policies. While we provide narrative accounts of what the measurements may say about the community, our real contribution is to *bracket* collaborative activity: to provide measurements that lend plausibility to certain interpretations of social dynamics and cast doubt on others. We proffer five investigations of the relationship between the policy environment and Wikipedia practices:

Policy citations track community concerns. We examine the evolution and continuity of the policy environment at the level of individual policies. With respect to a few known external events, adoption patterns suggest that citation patterns reflect community concerns.

Policy citations indicate salient work dimensions. We code policies for the social activity that the policy signals. We find that citations examined at this aggregate level illustrate salient work dimensions.

³In our analysis, all time series are given at a week granularity.

⁴This trend has also been observed for mainspace article edits (Kittur *et al.* 2007a).

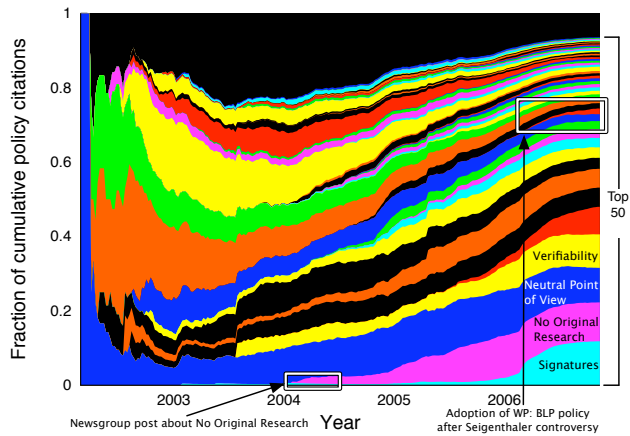


Figure 4: Each striation is a single policy. The width represents the fraction of cumulative policy citations. The 50 bottom striations are ordered from most cited policies in the last week of the dataset (bottom) to least cited (top). The top striation aggregates the remaining 142 policies.

Policy citation practices are converging. The community’s use of policies is gradually stabilizing. This holds within and between user classes.

Policy citations can mark transformations in participation. Enacting policy is one indication of participation in self-governance. We find that policies are being enacted to a large extent by users new to the practice. New citers also tend to cite the most popular policies, revealing the efficacy of socialization mechanisms.

The practice of citing policies is becoming increasingly commonplace. We examine the relationship between the two methods of invoking a policy – policy citation and keyword reference. While keyword references have historically been the dominant method, policy citations have become prevalent. This supports our claim that policy citations, a simple feature to extract, are a reliable social indicator.

Evolution of the policy environment

We characterize the development of the policy environment by first considering the relative distribution of policy citations amongst the individual policies over time. In Figure 4, each striation represents a single policy and the width represents the relative fraction of the cumulative policy citations. The figure demonstrates that the number of enacted policies is proliferating, yet there is continuity in their popularity. There is also consolidation in the most popular policies, with the top four accounting for 40% of all citations.⁵

The development of the policy environment is influenced by external events and responds to the needs of the community (Forte & Bruckman 2008). Both of these are evident in our trend data. First, we can see some policies

⁵We suspect that an imitation mechanism facilitated by the hyperlink mechanism underlies the generation of these citation patterns, similar to tag adoption in collaborative tagging systems (Golder & Huberman 2006).

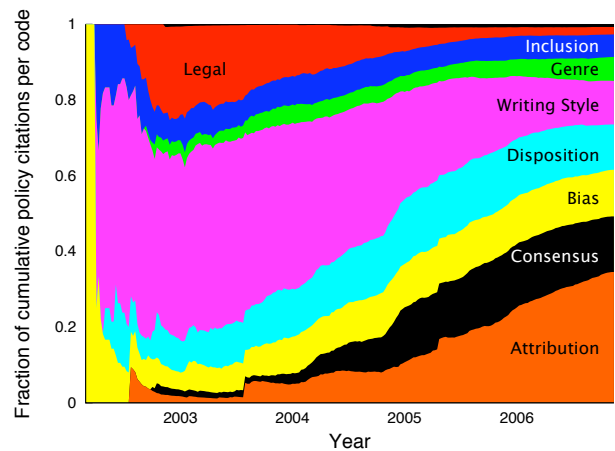


Figure 5: Each striation in this graph represents a single coded category. The width of the striation represents the fraction of cumulative policy citations for policies in the category. Striations are ordered from bottom to top according to their rank in the last week of the dataset.

“stick” because they manage to clarify key aspects of collaborative activity. The establishment and rapid growth of the No Original Research policy in early 2004 is indicative. The idea of “No Original Research” was first articulated in September 2003 through a newsgroup post by co-founder Jimmy Wales:

*If your viewpoint is held by an extremely small minority, then whether it’s true or not, whether you can prove it or not, it doesn’t belong in Wikipedia, except perhaps in some ancilliary article. Wikipedia is not the place for original research.*⁶

The concept resonated with the community, was codified and has grown to be the second most cited policy.

Wikipedia is not isolated from the rest of the world: institutional factors can affect policy enactment. For example, the establishment of the Biographies of Living Persons policy was largely in response to the controversy over John Seigenthaler Sr.’s article, when the media widely accused Wikipedia of harboring libelous information (Seigenthaler 2005). In Figure 4, we can see rapid adoption of this policy by the community in the aftermath.

Policy use reflects work patterns

Different policies speak to different principles of collaborative editing. In order to understand the relative importance of these principles, we categorized the policies by the type of social activity that a given policy citation may implicate. We employed a grounded approach (Strauss & Corbin 1990): first we read all the policy pages and developed a rough scheme. We then refined our categories and coded all the policies. Figure 6 gives our resulting codebook.

Figure 5 plots fractions of cumulative policy citations per code over time. It demonstrates a major shift in emphasis

⁶<http://lists.wikimedia.org/pipermail/wikien-l/2003-September/006715.html>

Code	Social Signal	Example
attribution	citing references	RS
consensus	consensus-seeking process	OWN
bias	neutrality of content or organization	NPOV
disposition	legitimacy of user actions or intent	CIVIL
writing style	low level style and organization	MOS
genre	defines Wikipedia	NOT
inclusion	non-legal issues of including content	NOTABLE
legal	legality of content or user actions	LEGAL

Figure 6: Codebook for classifying policies.

towards attribution work – making sure that Wikipedia is a tertiary source of knowledge where facts are attributed to reliable sources. This indicates a change in concern regarding content quality (Viégas, Wattenberg, & McKeon 2007; Forte & Bruckman 2008), which may also be a response to an increasing amount of negative press about reliability. As policies are often mobilized in heated conflicts, the trend may also signal a shift in the discursive tactics used to control article content from arguments about the text to strategies that question the source of information. We also see an increasing prominence of policies that implicate the consensus process and user behavior, speaking to the importance of policy use in mediating mass participation (Kriplean *et al.* 2007). In sum, policy citations track global shifts in discursive and attribution work, suggesting that policies may be valuable as micro-level indicators of work activity.

An interesting work dimension not accounted for in our coding scheme are those policies which are primarily enacted by the community to make up for lack of technical support. This is most clear with the *Signatures* policy. This policy explains the importance of “signing” discussion posts with four tildes, which the wiki software automatically substitutes with the user’s name and date of the post. Such a policy has arisen in the absence of proper technical support for discussions. Figure 4 shows that this policy has become the *most* cited policy in Wikipedia. The policy’s prominence is due to the incorporation of the ability to automatically rewrite unsigned discussions to include a reference to the *Signatures* policy into a widely used tool called the “AutoWikiBrowser” (AWB). This genre of work touches on an interesting design point: to what extent should enforcement of policy be encoded into the software? The need for the *Signatures* policy may be easily eliminated if the software were to automatically include a user’s signature when the edit occurred as part of a discussions. A more difficult case is the *Three-Revert Rule*. Although it would be amenable to technical enforcement, there are numerous cases where such strictness would be disruptive.

Stabilization of the policy environment

Over time, we might expect the enactment of the policy environment to stabilize, reflecting increasing diffusion of norms within the community. To study stability, we employ a measure of similarity that compares distributions of policies. We define a *policy vector* as a tuple of policy citation counts, with each tuple element corresponding to one policy. For example, if the policy environment had three policies, a policy vector may be {NPOV:4,RS:3,NOR:5}, where the *Neutral*

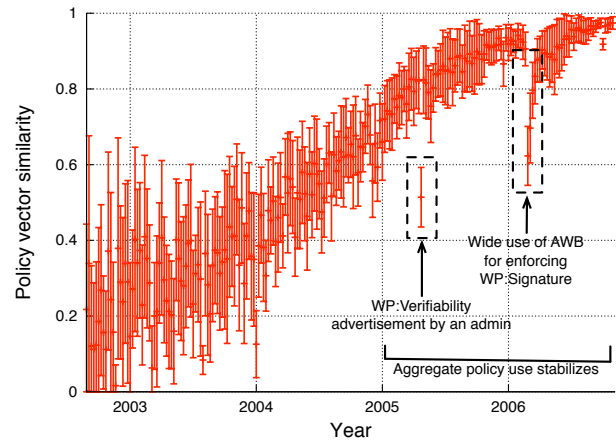


Figure 7: The convergence of policy vector similarities in a half-year moving window shows that the relative proportions of aggregate number of policy citations are stabilizing across Wikipedia. The lower outlier is due to an admin who posted an advertisement of the *Verifiability* policy across many talk pages to encourage its adoption. Another set of outliers is due to a wide adoption of AWB, a software tool for automating simple editing tasks.

point of view policy had four citations. The vector may represent the distribution of policy citations by registered users, or characterize the citations made on a single page. To compare two policy vectors v_1 and v_2 , we define their *similarity* as $\frac{v_1 \times v_2}{\|v_1\| \times \|v_2\|}$. This value is the angle between the two vectors in multi-dimensional policy space.⁷ A value of 0 indicates that the two policy vectors are orthogonal (different), and a 1 indicates that the two vectors are linearly dependent (similar). Our similarity measures are in the interval [0,1] because citations are non-negative. Note that this similarity analysis disregards the magnitude of policy citations; it compares the *ratios* of the magnitudes of policy citations.

We use similarity measures to compare the policy environment at different points in time. Figure 7 plots the mean and the standard deviation of a time-varying set of similarities. For each week, a set of similarities is calculated. The set contains the similarity values between that week’s policy vector and each of the policy vectors for the 26 previous weeks. Figure 7 shows that although the policy environment is expanding in policy quantity, as well as unique policy citers, it exhibits increasing stability.

Earlier we pointed out that admins cite more policy per user than registered users. But do these groups differ in the relative proportions of the types of policies they cite? Figure 8 plots the similarity between policy vectors for admin and registered users (*cross-group* stability). The figure also plots the similarity of policy vectors 4 weeks apart within each of the two user groups (*intra-group* stability). All measures of stability are converging to the *same policy vector*.⁸

⁷Taking the cosine of this angle would give us the more traditional cosine similarity metric.

⁸Examining similarity of policies employed by users partici-

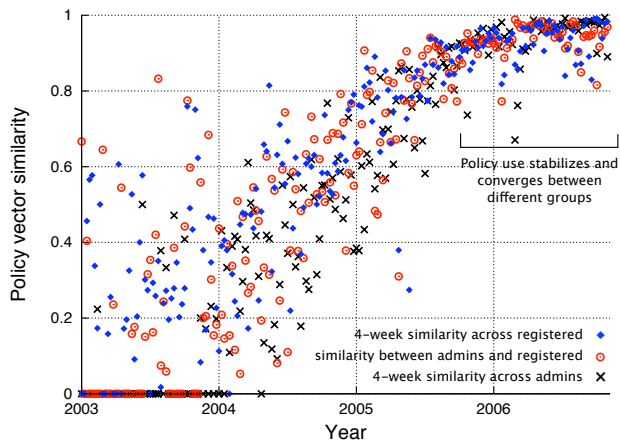


Figure 8: Policy vector similarities for admin and registered users are converging to 1, indicating that the two user groups are becoming similar in their policy citations. The convergence of similarities of policy vectors 4 weeks apart within each population indicates stability within user classes.

Policy use by veterans and first-time citers

Citing policy is a simple mechanism by which anyone can participate in Wikipedia’s governance structure. Here we examine how policy citations are distributed amongst contributors who have cited policy before (*veterans*) and those who have not. Figure 9 indicates that participation is inclusionary. It shows that in every week over 10% of the citers are first-timers.⁹ This may provide evidence of the accessibility of the policy environment by showing that policy citations are not exclusively made by veteran users. We think that the ease of hyperlinking to policy and the importance of socializing new users via the policy environment explains the prominence of users who have never cited policy.

We also compare the population of veteran policy citers and first-time citers via the types of policies that they cite. Figure 10 shows that users who cite policy for the first time on average begin with a top 5 policy in about 45% of the cases, matching the veteran citers. Comparable trends were observed when varying the top k policies. The similarity in fractions for first-time and veteran policy citers indicates a strong socialization of users in Wikipedia.¹⁰ The dip in Figure 10 by the veteran citers may be due to the proliferation of policies, while the more recent stabilization of the policy environment is likely responsible for the later steady increase in consolidation for both user groups.

pating in different Wikiprojects would be interesting given Forte *et al.*’s (2008) finding that WikiProjects are starting to establish their own local content rules.

⁹The curves for veteran and first-time policy citers are slowly diverging over time as the veteran user population comes to include users who have started citing policy regularly.

¹⁰Once again, this may signal an imitation mechanism at play, where new policy citers enact the citation patterns of veterans through observations of veterans.

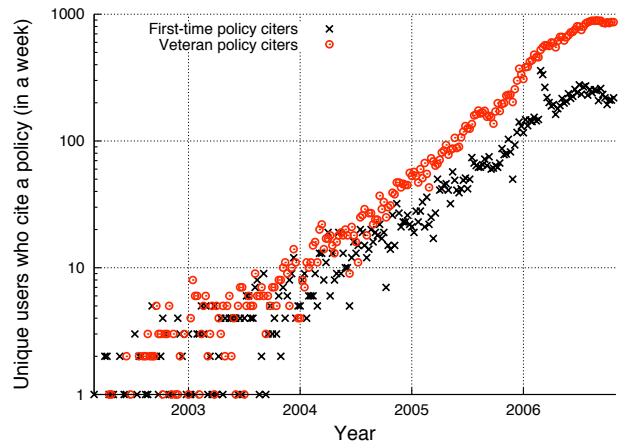


Figure 9: On average, unique users who cite policy for the first time account for over 10% of all unique users who cite a policy in a week. The high percentage indicates that participation in Wikipedia’s governance structure is inclusionary.

Policy references and communal vocabulary

Policy citations are not the only method of referencing policy – members may also reference policies in clear text (*keywords*) as part of a shared communal vocabulary.¹¹ We extracted common keyword references from talk pages for a few policies to characterize the enactment of the policy environment via keyword usage.¹² For each of these policies, we looked at the ratio of the hyperlinked policy citations to all references (citations + keyword references).

Figure 11 shows that policies are referenced by keyword more frequently than they are hyperlinked. We expected, however, that as policies migrate into the communal vocabulary the ratio would *decrease* over time. Instead, we see that linking to policy is becoming increasingly common. There are a couple explanations. First, hyperlinking to a policy is itself a norm that appears to be growing more prevalent. This indicates that the policy citation dataset is becoming more reflective of practice. Second, the Wikipedian population is exponentially growing and links are an easy shorthand socializing mechanism (Viégas *et al.* 2007). Interestingly, the figure shows that different policies have different temporal relationships in enactment style. For example, the oldest of the seven policies in the figure – the Neutral Point Of View Policy – is primarily referenced through keywords. The relatively recent Reliable Sources guideline, on the other hand, is hyperlinked in over 50% of total references.

Discussion

Studying social practice quantitatively. Quantitative studies of Wikipedia have generally taken as their main unit

¹¹Our measures of policy citations are therefore a lower bound of policy references in discussions.

¹²We did not attempt to mine all possible keyword references. Consider the no original research policy that is typically referenced by the ‘OR’ keyword. Disambiguating the use of ‘or’ as a disjunctive from its use as a keyword reference is difficult.

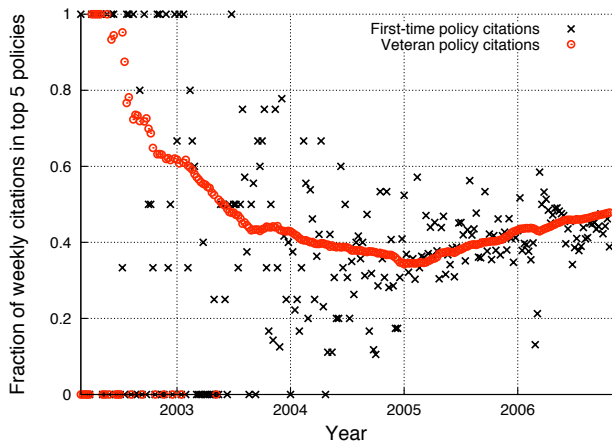


Figure 10: On average, users who cite a policy for the first time cite a top 5 policy in 45% of the cases, closely matching users who have previously cited a policy (veterans).

of analysis the edit to a page, text persistence, and/or reverts to pages. Undoubtedly, these metrics can be used to generate complex and interesting models. (*c.f.* Adler *et al.*'s (2007) reputation model). However, they are limited in their ability to bracket complex social behavior. Priedhorsky *et al.* (2007) recognized a similar issue in the “author-based” metrics typically used to study Wikipedia and instead advocate for “reader-based” data sources. Based on this idea, they synthesized page view data from a number of sources to create a model of the impact of an edit.

Page view data, however, is relevant primarily for modeling main article space. Metrics such as policy citations are more appropriate for studying the conditions of Wikipedia's production. We suggest that researchers might focus on mining semantically rich indicators of social activity such as page locking, requests for comments, user blocking, WikiProjects, userboxes, barnstars, and many others in order to better understand and extend the affordances of Wiki technology for supporting large-scale collaboration. For example, policy citations may be used as a work-aware indicator of editing hot-spots, providing a new source of data for conflict visualizers (Kittur *et al.* 2007b) and awareness mechanisms (Suh *et al.* 2008).

There are significant challenges for interpreting trends in these datasets. Temporal trends are particularly difficult, as the social practices that these datasets reflect are themselves evolving. Consider policy citations: it is difficult to understand the ramifications of a particular citation without first understanding the version of the policy environment that it references. Moreover, as demonstrated earlier, the practice of citing policy is itself being established; just because policy citations were rarer in 2005 than in 2006 does not necessarily mean that the policy environment was less important. We are interested in principled techniques that have been brought to bear on studying developing social practices.

Another challenge for studying social practice in Wikipedia is to correlate activity occurring in different social spaces. Discussions on the talk pages take place as the arti-

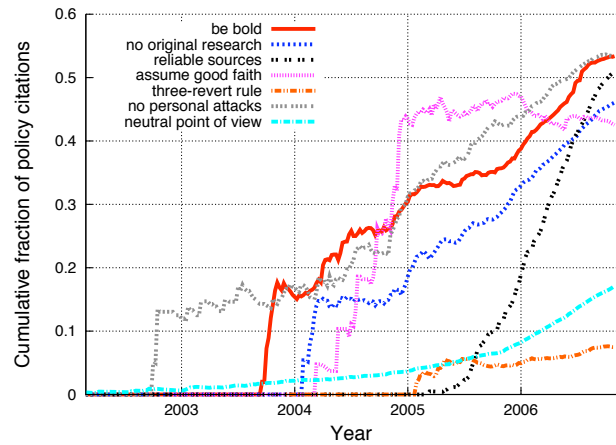


Figure 11: Policy citations and keyword references can both invoke a policy. In general, the ratio of policy citations to the total number of references to policy (citations and keyword references) is increasing over time. This indicates that linking to policy is becoming common place even though keyword references continue to dominate.

cle is edited, users often simultaneously leave messages on each other's user talk pages, conflicts are often escalated to forums dedicated to informal and formal dispute resolution, and groups of editors formed around WikiProjects initiate collective action. Techniques that can extract related activities across these spaces could serve to elucidate higher level collaborative patterns. Such models may facilitate tools that help Wikipedians navigate to events of interest and support Wikipedians as they react to these events.

Design of online communities. The policy environment – user editable, reflective of practice, and easily citable – is an integral mechanism for maintaining quality and organization in Wikipedia. We believe that in the future we will see other participatory communities collaborating in high-stakes domains that may benefit from such mechanisms (*e.g.* public deliberation (Borning *et al.* 2008)). Our study gives evidence that providing a means for a community to express shared strategies of action, codes of conduct, and content norms is a powerful and inclusionary way to enable self-governance, following from early work in CSCW that highlights the necessity of social spaces for articulation work (Schmidt & Bannon 1992).

As the product of many years of large-scale effort, the policy environment captures experiential knowledge about consensus-based collective action. A potentially fruitful area for future investigation is the transplantability of the policy environment itself: to what extent and under what conditions can new collaborative endeavors be seeded with Wikipedian policies?

Conclusion

Our research motivates the use of policy citations to study the establishment of practice in Wikipedia. We took the view that we can learn about the relative importance of differ-

ent policies through their *usage*, not their formal definition and development on the Wikipedia policy pages. We have demonstrated that policy citations are situated at an interesting position between normative and formal structure and that tracking them gives insight into core aspects of Wikipedia's conditions of production. We invite other researchers to join us in studying policies as an interesting social index.

Acknowledgements

We would like to thank Eytan Adar for very helpful comments. Kriplean would also like to thank HP Lab's Information Dynamics team, especially Scott Golder and Bernardo Huberman, for their support of the initial investigation of policy citations during the summer of 2006. We are also indebted to Wikimedia for providing public data dumps.

References

- Adler, T. B., and de Alfaro, L. 2007. A Content-driven Reputation System for the Wikipedia. In *Proceedings of WWW*, 261–270.
- Almeida, R. B.; Mozafari, B.; and Cho, J. 2007. On the Evolution of Wikipedia. In *Proceedings of ICWSM*.
- Arguello, J.; Butler, B. S.; Joyce, E.; Kraut, R.; Ling, K. S.; Rosé, C.; and Wang, X. 2006. Talk to me: Foundations for Successful Individual-group Interactions in Online Communities. In *Proceedings of CHI*, 959–968.
- Borning, A.; Friedman, B.; Davis, J.; Gill, B.; Kahn, P.; Kriplean, T.; and Lin, P. 2008. Laying the Foundations for Public Participation and Value Advocacy: Interaction Design for a Large Scale Urban Simulation. In *Proceedings of DGO*, to appear.
- Bryant, S. L.; Forte, A.; and Bruckman, A. 2005. Becoming Wikipedian: Transformation of Participation in a Collaborative Online Encyclopedia. In *Proceedings of GROUP*, 1–10.
- Buriol, L.; Castillo, C.; Donato, D.; Leonardi, S.; and Millozzi, S. 2006. Temporal Evolution of the Wikigraph. In *Proceedings of Web Intelligence*, 45–51.
- Cosley, D.; Frankowski, D.; Terveen, L.; and Riedl, J. 2007. SuggestBot: Using Intelligent Task Routing to Help People Find Work in Wikipedia. In *Proceedings of IUI*, 32–41.
- Fisher, D.; Smith, M.; and Welser, H. T. 2006. You Are Who You Talk To: Detecting Roles in Usenet Newsgroups. In *Proceedings of HICCS*, 59.2. Washington, DC, USA: IEEE Computer Society.
- Forte, A., and Bruckman, A. 2008. Scaling Consensus: Increasing Decentralization in Wikipedia Governance. *Proceedings of HICCS* 157–166.
- Golder, S., and Huberman, B. A. 2006. Usage Patterns of Collaborative Tagging Systems. *Journal of Information Science* 32(2):198–208.
- Kittur, A.; Chi, E.; Pendleton, B. A.; Suh, B.; and Mytkowicz, T. 2007a. Power of the Few vs. Wisdom of the Crowd: Wikipedia and the Rise of the Bourgeoisie. In *Alt.CHI*.
- Kittur, A.; Suh, B.; Pendleton, B.; and Chi, E. 2007b. He says, She says: Conflict and Coordination in Wikipedia. In *Proceedings of CHI*, 453–462.
- Kollock, P., and Smith, M. 1996. *Managing the Virtual Commons: Cooperation and Conflict in Computer Communities*. Jahn Benjamins. 109–128.
- Kriplean, T.; Beschastnikh, I.; McDonald, D. W.; and Golder, S. 2007. Community, Consensus, Coercion, Control: CS*W or How Policy Mediates Mass Participation. In *Proceedings of GROUP*, 167–176.
- Ortega, F., and Barahona, J. M. G. 2007. Quantitative Analysis of the Wikipedia Community of Users. In *Proceedings of WikiSym*, 75–86.
- Pentzold, C., and Seidenglanz, S. 2006. Foucault@Wiki: First Steps Towards a Conceptual Framework for the Analysis of Wiki Discourses. In *Proceedings of WikiSym*, 59–68.
- Priedhorsky, R.; Chen, J.; Lam, S.; Panciera, K.; Terveen, L.; and Riedl, J. 2007. Creating, Destroying, and Restoring Value in Wikipedia. In *Proceedings of GROUP*, 259–268.
- Schmidt, K., and Bannon, L. 1992. Taking CSCW Seriously. *Computer Supported Cooperative Work* 1(1):7–40.
- Seigenthaler, J. 2005. A False Wikipedia 'Biography'. *U.S.A. Today*.
- Strauss, A., and Corbin, J. 1990. *Basics of Qualitative Research*. Sage.
- Suh, B.; Chi, E. H.; Kittur, A.; and Pendleton, B. A. 2008. Lifting the Veil: Improving Accountability and Social Transparency in Wikipedia with WikiDashboard. In *Proceedings of CHI*, to appear.
- Turner, T. C.; Smith, M. A.; Fisher, D.; and Welser, H. T. 2005. Picturing Usenet: Mapping Computer-Mediated Collective Action. *Journal of Computer-Mediated Communication* 10(4).
- Viégas, F. B.; Wattenberg, M.; Kriss, J.; and van Ham, F. 2007. Talk Before You Type: Coordination in Wikipedia. *Proceedings of HICSS* 78–78.
- Viégas, F. B.; Wattenberg, M.; and Dave, K. 2004. Studying Cooperation and Conflict Between Authors with History Flow Visualizations. In *Proceedings of CHI*, 575–582.
- Viégas, F. B.; Wattenberg, M.; and McKeon, M. M. 2007. The Hidden Order of Wikipedia. In *HCI (15)*, 445–454.
- Voss, J. 2005. Measuring wikipedia. In Ingwersen, P., and Larsen, B., eds., *Proceedings of ISSI*, volume 1, 221–231.
- Wattenberg, M.; Viégas, F.; and Hollenbach, K. 2007. Visualizing Activity on Wikipedia with Chromograms. In *Proceedings of INTERACT*, 272–287.
- Whittaker, S.; Terveen, L.; Hill, W.; and Cherny, L. 1998. The Dynamics of Mass Interaction. In *Proceedings of CSCW*, 257–264. New York, NY, USA: ACM Press.
- Wilkinson, D. M., and Huberman, B. A. 2007. Cooperation and Quality in Wikipedia. In *Proceedings of WikiSym*, 157–164.
- Wu, F., and Weld, D. S. 2007. Autonomously Semantifying Wikipedia. In *Proceedings of CIKM*, 41–50.