

American Behavioral Scientist

<http://abs.sagepub.com/>

Hot Off the Wiki: Structures and Dynamics of Wikipedia's Coverage of Breaking News Events

Brian Keegan, Darren Gergle and Noshir Contractor

American Behavioral Scientist published online 10 January 2013

DOI: 10.1177/0002764212469367

The online version of this article can be found at:

<http://abs.sagepub.com/content/early/2013/01/09/0002764212469367>

Published by:



<http://www.sagepublications.com>

Additional services and information for *American Behavioral Scientist* can be found at:

Email Alerts: <http://abs.sagepub.com/cgi/alerts>

Subscriptions: <http://abs.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

>> [OnlineFirst Version of Record](#) - Jan 10, 2013

[What is This?](#)

Hot Off the Wiki: Structures and Dynamics of Wikipedia's Coverage of Breaking News Events

American Behavioral Scientist

XX(X) 1–28

© 2013 SAGE Publications

Reprints and permission: <http://www.sagepub.com/journalsPermissions.nav>

DOI: 10.1177/0002764212469367

<http://abs.sagepub.com>



Brian Keegan¹, Darren Gergle¹, and
Noshir Contractor¹

Abstract

Wikipedia's coverage of breaking news and current events dominates editor contributions and reader attention in any given month. Collaborators on breaking news articles rapidly synthesize content to produce timely information in spite of steep coordination demands. Wikipedia's coverage of breaking news events thus presents a case to test theories about how open collaborations coordinate complex, time-sensitive, and knowledge-intensive work in the absence of central authority, stable membership, clear roles, or reliable information. Using the revision history from Wikipedia articles about over 3,000 breaking news events, we investigate the structure of interactions between editors and articles. Because breaking article collaborations unfold more rapidly and involve more editors than most Wikipedia articles, they potentially regenerate prior forms of organizing. We analyze whether the structures of breaking and nonbreaking article networks are (a) similarly structured over time, (b) exhibit features of organizational regeneration, and (c) have similar collaboration dynamics over time. Breaking and nonbreaking article exhibit similarities in their structural characteristics over the long run, and there is less evidence of organizational regeneration on breaking articles than nonbreaking articles. However, breaking articles emerge into well-connected collaborations more rapidly than nonbreaking articles, suggesting early contributors play a crucial role in supporting these high-tempo collaborations.

Keywords

Wikipedia, news, collaboration, social network

¹Northwestern University, Evanston, IL, USA

Corresponding Author:

Brian Keegan, Northwestern University, Evanston, 2240, Campus Drive, Evanston, IL 60208, USA

Email: bkeegan@northwestern.edu

In the wake of disasters, scandals, and other unexpected news events, unique forms of social behavior and organization emerge to support information dissemination, response, and sense making (Quarantelli & Dynes, 1977). Social media spaces exhibit analogous processes as citizens and responders use mobile phones, text messaging, GPS, and microblogging to organize ad hoc responses, disseminate information, and provide social support (Shklovski, Palen, & Sutton, 2008). Given Wikipedia's prominence as a popular reference website, it is an obvious locus for information seeking and sensemaking following unexpected and highly salient events. However, developing an authoritative account on an inherently collaborative project like Wikipedia where "anyone can edit" is a particularly complicated task. Editors have diverse motivations and access to resources, many have never worked together before and may never work together again, and these editors have substantial volition to participate as much or as little as they prefer. In this volatile information environment no central authority assigns tasks, makes decisions, or enforces rules. Extant organizational theory would lead us to believe these factors would inhibit effective collaboration among editors.

Nevertheless, practice proves otherwise. Like other sociotechnical systems, such as mobile phones and microblogs, wikis appear to be well suited to supporting the temporary organizations that emerge to document breaking news events, and have been cited as exemplars of timeliness, breadth, and reliability in the wake of disasters such as the 2007 Virginia Tech massacre (Cohen, 2007). In an internal 2009 memo, the Associated Press notes,

[A] new pattern of [news] consumption was validated in the confusing minutes that followed the first reports of [Michael Jackson's death]: users shared; they searched and they clicked on Wikipedia. . . . The new routine of Twitter-to-Google-to-Wikipedia contrasts sharply with the behavior of users [after Princess Diana's death]. . . . The key to Wikipedia's success is that its pages are designed to catch traffic, provide key information and then send users on their way to deeper engagement on the subjects they're interested in. . . . For large portions of the Internet audience, [Wikipedia was] an authoritative source working to verify an important news development. (Seward, 2009)

Since 2003, the top 25 Wikipedia articles with the most contributors every month consist nearly exclusively of articles pertinent to current events (Wikipedia, 2011). For example, the articles that attracted the most contributors in February 2011 included "2011 Egyptian Revolution," "Super Bowl XLV," and "2011 Christchurch Earthquake" (Wikipedia, 2011), and similar results are found regarding the number of unique edits. Moreover, Wikipedia page views in any given week or month likewise demonstrate a substantial bias toward articles about current events (Mituzas, 2011), suggesting Wikipedia editors are engaged in a unique type of "citizen journalism" (Deuze, Bruns, & Neuberger, 2007).

This article examines Wikipedia articles about natural disasters, technological accidents, and violent conflict across 20 years to understand how emergent response groups in high-tempo and online collaborative contexts are similarly structured over time, exhibit features of organizational regeneration, and have distinct evolutionary dynamics from nonbreaking articles. Although the growth of activity on breaking articles contrasts with the stability of nonbreaking articles, both types of articles exhibit surprising similarity in their structures. However, breaking articles produce large, interrelated collaborations immediately following their creation, whereas nonbreaking articles take a year or more to exhibit similar connectivity. Our findings about this unique class of Wikipedia articles suggest alternative approaches to theorizing about large-scale online collaborations.

Theoretical Background

Studies of peer production in sociotechnical systems such as online communities often presuppose that membership in and motivations to contribute to these communities are relatively stable and constant, yet such communities often self-organize and coordinate their activities under conditions of unstable and sudden collective action. In this article we align with the “crisis informatics” approach, which analyzes how citizens and responders employ sociotechnical systems such as mobile phones, wikis, and microblogging to organize ad hoc responses, process and disseminate information and provide social support (Shklovski et al., 2008; Starbird, Palen, Hughes, & Vieweg, 2010). Sociotechnical systems like Twitter, Facebook, and Wikipedia play an important role in disseminating information about breaking news events such as the Tōhoku tsunami or the death of Osama bin Laden, yet our understanding of how online communities rapidly self-organize to support this knowledge work remains largely anecdotal and theoretically unmoored.

High-tempo contexts are characterized by nonroutine and urgent work, abrupt consequences, intense attention, and ephemeral teams. Coordination in volatile environments such as disaster response or emergency medicine demands high levels of heedful and interrelated action, knowledge integration, and information processing (S. Brown & Eisenhardt, 1997). Some organizations such as emergency rooms and aircraft carrier crews respond to demands for high tempo and high reliability action by defining clear group membership, tasks, and expertise (Faraj & Xiao, 2006; Weick & Roberts, 1993). These temporary organizations are governed through networks of relations rather than lines of authority, which leads to coordination mechanisms emphasizing reciprocity, socialization, and reputation (Jones, Hesterly, & Borgatti, 1997). However, in contexts such as disaster response, groups lack preexisting role structures and reputations and are unable to routinize practices or rely on prior expertise. These *emergent response groups* are unique because members have diverse motivations, mixed perspectives, varied resources to contribute, and can come and go as they please. Editors of Wikipedia’s breaking news articles face constraints analogous to these emergent

response groups. Following a breaking news event such as an earthquake or commercial airliner crash, the facts must be reconstructed, negotiated, and integrated into the account even as new information continues to unfold. Editors self-organize and draw on diverse skills, expertise, and motivations to fulfill particular social roles that support collaboration (Welser et al., 2011). Some of these contributors have previously collaborated on a breaking news event, whereas others may make several contributions and never edit Wikipedia again. The responsibilities for integrating and updating content, reverting vandalism, formatting citations, and mediating disputes are likewise diffused among the members of the collaboration.

Because breaking news articles are coauthored in high-tempo contexts that are rarely found on “typical” Wikipedia articles, we expect breaking articles’ collaborations to exhibit structural features that are distinct from other Wikipedia collaborations. In many online collaborations and communities, the majority of activity is centralized in a handful of users (Wilkinson, 2008). Teams operating under conditions of uncertainty and complexity are more effective when employing decentralized interaction patterns as this promotes more contributions and reduces the risk of overload for central integrators. However, teams operating under conditions of threat and time pressure are more effective when employing centralized interaction patterns as this promotes greater control and simpler information processing schemas. Prior work suggests the interactions between these constraints and their influences on team performance are complicated (Argote, Turner, & Fichman, 1989; T. Brown & Miller, 2000). Unlike typical Wikipedia articles, breaking news articles place demands on users to collaborate in contexts that are simultaneously uncertain as well as time sensitive, which demands editors adopt implicit coordination mechanisms that lead to either greater concentration or distribution of work across editors (Kittur & Kraut, 2008; Kittur, Lee, & Kraut, 2009). Because of the differences in the coordination demands of breaking news articles versus typical Wikipedia articles, we expect that (a) the concentration and distribution of editing activity on Wikipedia’s breaking articles will differ significantly from typical articles and (b) this concentration and distribution of work will be consistent across events occurring in different years.

Hypothesis 1a: Breaking news articles will exhibit concentrations of editing activity that are significantly different from nonbreaking articles’ concentrations.

Hypothesis 1b: Breaking news articles’ concentrations of editing activity will be consistent for events occurring in different years.

Organizational theory has grappled with coordination and self-organization in temporary teams by examining the extent to which participants regenerate, adapt, and improvise roles and routines used in previous projects and collaborations (Bechky, 2006; Birnholtz, Cohen, & Hoch, 2007). Because the reenactment of structures among a cohort of regular participants may contribute to high performance of the collective, coauthorship of breaking articles on Wikipedia may not be one-off but rather involve

editors who have repeatedly worked together or even specialize in editing content about breaking articles. Wikipedia's ability to cover breaking news events may reflect the regeneration of organizational forms capable of coordinating high-tempo work through the self-organization of editors who have previously participated in breaking news collaborations. The presence of some of the same editors across collaborations suggests they have greater capacity for knowledge coordination such as group mind to manage routines and delegate tasks, credibility to manage conflicts, and expertise to adapt processes and information in high-tempo work. If breaking article collaborations adapt to task demands by relying on contributions from editors who have previously contributed to other breaking articles, they should exhibit a greater tendency for repeated participation among editors for incidents occurring in different years.

Hypothesis 2: Breaking news articles will have more editors who previously contributed to other breaking news articles than nonbreaking news articles' editors contributing to other nonbreaking articles.

Furthermore, the structure and regeneration of these collaborations may change over time. Understanding these dynamics can provide insight into how large and connected networks balance access to diverse and novel information with the costs of reduced information flow and heightened coordination demands (Aral & Alstynne, 2011). In large and connected networks, some individuals have highly central positions, others inhabit highly embedded and clustered neighborhoods, and still others are boundary spanners who broker connections. Each of these positions represents individuals with special resources, strong ties for action, and connections across diverse groups, and all are prerequisites for high levels of organizational performance (Burt, 2004). The success of organizations operating in high-tempo contexts necessarily depends on their ability to rapidly assemble resources and create the network of communication and resource-sharing ties needed to respond to the event (Faraj & Xiao, 2006). The coordination demands that give rise to the concentration and distribution of activity on breaking news articles (Hypothesis 1) combined with the regeneration of these organizational memberships across collaborations (Hypothesis 2) suggests the networks of breaking news articles and their editors should exhibit different patterns of connectivity and shared coauthorship over time compared to nonbreaking or historical articles.

Hypothesis 3: The distribution of activity on breaking article collaborations will differ significantly from nonbreaking articles' distributions in their early histories.

Prior Work

As is the case with many online communities, the majority of contributions to Wikipedia come from a fraction of the entire user base (Kittur, Chi, Pendleton, Suh,

& Mytkowicz, 2007; Panciera, Halfaker, & Terveen, 2009). Previous studies of coordination in Wikipedia suggest implicit and informal coordination mechanisms can support article development when the article is young and intensively worked on by a dense cohort of authors. As an article ages, coordination shifts toward more explicit and formal mechanisms such as discussion (Kittur et al., 2009). Scholarship has only begun to examine high-tempo collaboration mechanisms following unexpected and surprising events and Wikipedia articles about current news events are sites of collective memory, sense making, and commemoration (Ferron & Massa, 2011; Keegan, 2011). Previous studies of the structures or dynamics of Wikipedia breaking news coauthorship have only examined a single category of articles limiting the generalizability of their findings (Keegan, Gergle, & Contractor, 2011, 2012). This study substantially expands the sample over prior studies by analyzing a corpus of 3,233 Wikipedia articles across a much wider range of genres for breaking and nonbreaking events since 2001 to compare the structure and dynamics of both breaking and nonbreaking articles.

Data and Method

Data

We identified seven broad categories of Wikipedia articles likely to include breaking news events. These include conflicts (e.g., wars, battles, political unrest), crimes (e.g., murders, kidnappings, and terrorism), fires (e.g., building fires, wildfires, and explosions), health disasters (e.g., disease outbreaks), industrial accidents (e.g., spills, mine collapses), natural disasters (e.g., hurricanes, earthquakes, and tornadoes), and transportation accidents (e.g., airplane crashes, train collisions, road accidents). These categories are also categorized by incident year. For example, the Tōhoku earthquake appears in the “2011 earthquakes” category because the incident itself occurred in 2011 even if the article was written at some subsequent time.

Using the English Wikipedia’s Application Programming Interface (<http://en.wikipedia.org/w/api.php>), we extracted revision histories for every article in these categories and their subcategories in January 2012. Some events, such as the wars in Iraq and Afghanistan and the 2010 Deepwater Horizon explosion and spill, had dedicated subcategories containing dozens of related articles. These subcategories included articles and lists about people, places, and events of relatively minor importance or limited similarity to the parent event itself (e.g., political leaders, recording artists releasing benefit CDs, nonprofit organizations); these extraneous subarticles were manually identified and removed to yield 3,233 articles focused on the events and incidents alone. These articles represent 195,831 unique editor–article interactions from 114,153 unique users from September 2001 to January 2012.

Revision history data include editor name and ID (or IP address), article name and ID, and time stamp. Based on these data, we extracted article-level attributes such as

Table 1. Examples of Breaking News Articles From Each Year

2001	Enron, American Airlines Flight 587
2002	U.K. firefighter dispute, 2002–2003
2003	Space Shuttle Columbia disaster, Northeast blackout of 2003
2004	Second Battle of Fallujah, 2004 Madrid train bombings
2005	July 7, 2005, London bombings, Hurricane Katrina
2006	2006 Israel–Gaza conflict, Comair Flight 191
2007	2007 U.K. floods, Writers Guild of America strike
2008	2008 Tibetan unrest, 2008 South Ossetia War
2009	2009 flu pandemic, Air France Flight 447
2010	Deepwater Horizon explosion, 2010 Haiti earthquake
2011	2011 Mumbai bombings, Tōhoku earthquake and tsunami

whether or not the article is breaking or nonbreaking (described below). To construct coauthorship networks, the revision histories for articles in a given year are converted to weighted bipartite edge lists of editor ID and article ID. A single edge represents the number of times a single Wikipedia editor made contributions to a single article. We summarize basic descriptive networks in Table 2.

Breaking News Classification

To identify structural differences between high-tempo collaborations around breaking news events and typical Wikipedia collaborations, we identify three classes of articles: breaking articles, nonbreaking articles, and historical articles. Breaking and nonbreaking articles are about events that are contemporaneous with Wikipedia's existence since January 2001, whereas historical articles are about events between January 1990 and January 2001. These breaking and nonbreaking categories differentiate the temporal proximity between the article's creation and the date of the incident itself. Breaking articles were identified by examining the lag between an event (or the end of an event in the case of an ongoing situation such as a battle) and the creation date of its corresponding article. Examples of breaking news articles are given in Table 1.

Computing the difference between the date of the first edit and the date of the event itself, we observe the distribution of article creation lags plotted in Figure 1. Negative values of article creation lag are an artifact of the coding for noninstantaneous events in which the first edit to the article occurred before the incident ended. Taking a 1-day article creation lag to be the cutoff that differentiates breaking articles from nonbreaking articles, we observe 1,212 breaking articles and 2,074 nonbreaking articles. We plot the distribution of breaking, same-year, and nonbreaking articles in Figure 3.

Table 2. Summary Bipartite Network Statistics for Breaking and Nonbreaking Articles

	Nodes		Links		Density		Fraction of nodes in the largest connected component (LCC)		Clustering	
	Breaking	Nonbreaking	Breaking	Nonbreaking	Breaking	Nonbreaking	Breaking	Nonbreaking	Breaking	Nonbreaking
2001	2,046	7,415	2,072	9,863	0.506	0.000872	1.0	1.0	0.981	0.724
	2	154					1.0	1.0	0.0127	0.0305
2002	84	8,669	84	11,372	1.0	0.0155	1.0	1.0	1.0	0.752
	1	167					1.0	1.0	1.0	0.0443
2003	2,336	11,786	2,482	15,169	0.177	0.000674	1.0	1.0	0.9198	0.773
	6	173					1.0	1.0	0.0166	0.0363
2004	3,163	10,803	3,632	13,718	0.0717	0.000675	1.0	0.999	0.838	0.777
	16	175					1.0	0.989	0.0185	0.0314
2005	6,628	10,880	8,295	14,098	0.0278	0.000478	1.0	0.999	0.766	0.765
	45	195					1.0	0.989	0.0201	0.0371
2006	12,022	13,137	16,504	18,042	0.0129	0.00360	1.0	0.999	0.731	0.726
	240	155					1.0	0.995	0.0187	0.0292
2007	13,137	13,098	18,042	16,982	0.00886	0.00649	1.0	1.0	0.703	0.771
	155	238					1.0	1.0	0.0211	0.0339
2008	12,542	12,055	17,831	15,647	0.008831	0.00754	1.0	1.0	0.685	0.770
	171	240					1.0	1.0	0.0233	0.0308
2009	10,270	10,164	14,241	12,753	0.00726	0.00655	1.0	1.0	0.707	0.787
	191	209					1.0	1.0	0.0284	0.0354
2010	8,566	10,467	12,610	13,054	0.00751	0.00636	1.0	0.999	0.660	0.793
	196	195					1.0	0.995	0.0288	0.0294
2011	10,532	4,059	15,261	5,620	0.00766	0.00735	1.0	0.999	0.663	0.701
	189	173					1.0	0.994	0.0257	0.0553

For cells with two values, editor metrics are on top and article metrics are on bottom.




Figure 1. Distribution of article creation lags for articles with first revisions occurring the same year as the incident itself
Most articles' first revisions occur less than 1 day after the incident, which we then classify as "breaking news articles."

Coauthorship Network Evolution

To capture the evolution of these collaborations at the level of the article, we transform the edit history of each article to make the first edit to each article $t = 0$. Following this alignment, we assess the extent to which the coauthorship networks are similarly structured at the same relative time in each article's revision history. This process is illustrated in Figure 2, which contains example timelines for two breaking and two nonbreaking articles. Here, breaking articles have a greater concentration of activity in the early stages and nonbreaking articles have a more even distribution of activity throughout time. Because the distributions of editor revisions in breaking articles are strongly left skewed, the network is thresholded at 13 quasi-logarithmic time points. These thresholds correspond to 1 hour, 12 hours, 1 day, 2 days, 4 days, 1 week, 2 weeks, 1 month, 1 quarter, 1 year, 2 years, 4 years, and 1 decade after the first revision to an article. These thresholds are used to query the network structures at different points in time that we describe below.

We adopt two approaches to model changes in network structure over time. The first is a *cumulative network*, which aggregates all edits made by all editors to all articles before a given time threshold. In Figure 2, it would provide a representation of all previous contributions to a set of articles up to the third threshold line but exclude any revisions after the threshold. The second approach is a *snapshot network* that




Figure 2. Schematic representation of edit history in true time, and its transformation and alignment into relative time, for two breaking articles and two nonbreaking articles. Red dots reflect the first edits made to an article by an editor; black dots reflect subsequent revisions to that article. Examples of time thresholds are given by vertical gray dotted lines.

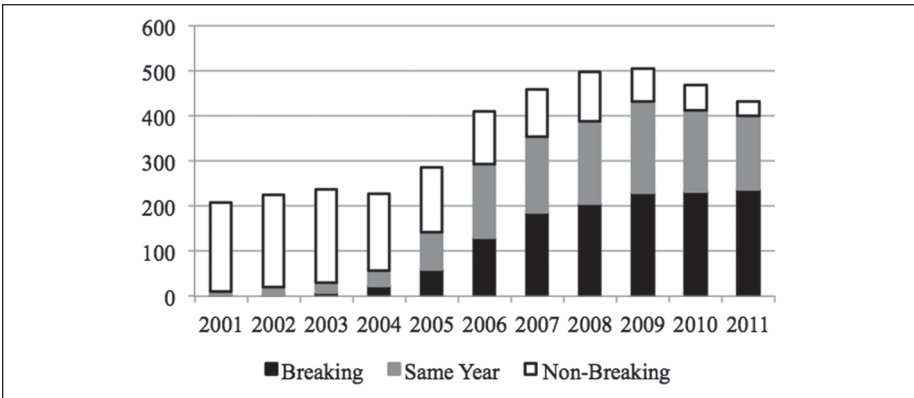


Figure 3. Frequency of article types per event year

looks only at edits that occurred *between* adjacent time thresholds. In Figure 2, it would include revisions made between the second and third threshold lines but exclude any other revisions that occur beforehand or afterward. Thus, the snapshot network provides a window into the structure of the collaboration shifts at different points throughout the history of an article category.

Results

A total of 64,272 unique editors had 82,254 distinct connections to 1,034 breaking articles and 61,571 unique editors had 113,577 distinct connections to 2,159 non-breaking articles. Figure 3 plots the frequency of different article types by event year. First, we observe a nearly monotonic increase in the absolute number of

articles from 207 articles in 2001 to 456 articles in 2011 with a peak of 527 in 2009. Nonbreaking articles about 2011 events are necessarily right censored as the data were collected in late January 2012. Although the number of disasters, accidents, and other catastrophic events certainly has not tripled around the globe in recent years, this article growth reflects both the increasing popularity of Wikipedia from 2001 through 2007 and the growth of articles about current events despite slowing article creation and editor participation after 2007. This trend points to the growing role Wikipedia plays as a global memory place for documenting, encoding, and commemorating collective memories of traumatic events (Ferron & Massa, 2011; Pentzold, 2009). Second, we observe a distinct shift in the relative distribution of breaking articles. In the years 2001 through 2005, breaking articles made up less than 25% of all event articles. In 2006, 48% of event articles were written as breaking articles, and in 2010 and 2011 that percentage exceeded 70%. Removing breaking and same-year articles, the trend remains: Nonbreaking articles decreased from 197 in 2001 (95%) to 41 in 2011 (8%). For the subsequent analysis, two groups were used where the first consisted of breaking articles and the second included both same-year and nonbreaking articles.

Structural Similarity

Descriptive network statistics for the entire network are summarized in Table 2. The number of nodes for breaking and nonbreaking articles reproduces the observations from Figure 3. There is a marked increase of contributions from unique editors on breaking articles, whereas nonbreaking articles exhibit more stable numbers of editors over time. Both of these factors influence how the density of the collaboration networks change over time: Breaking article collaborations become more sparse as they grow larger and nonbreaking articles remain stable. Breaking article collaborations were initially more dense than nonbreaking articles, suggesting less distribution of work, but collaborations around breaking articles in later years exhibit similar densities to nonbreaking articles. Every single revision to breaking articles ends up in a single giant component for each year, meaning these collaborations are all connected with each other through shared editorship, despite being about vastly different topics. In contrast, nonbreaking articles sometimes involve contributors editing articles that no one else edits in some years. Clustering coefficients for editors and articles capture the extent to which editors share articles in common or articles share editors in common, respectively (Latapy, Magnien, & Vecchio, 2008). The growth of breaking article editorship over time has resulted in a decrease in editor “familiarity” as they are less likely to share breaking articles in common in later years. Conversely, nonbreaking articles’ editors have stable clustering patterns over time, a feature that is further unpacked in Figure 6 below. Nonbreaking articles exhibit a greater tendency to have editors in common than breaking articles. The lack of clustering suggests breaking articles exhibit less distributed editing activity than nonbreaking articles.




Figure 4a. Editor degree distributions, by year for breaking articles

In Figure 4, we plot the degree and weight distributions for the editors and articles for events occurring in each year. In both breaking and nonbreaking articles, these distributions exhibit classic long-tail distributions. The fact that hundreds of editors edit only a single article whereas a handful of editors edit most or all of the articles is an example of *editor centralization*. Likewise, the observation that dozens of articles have only a few editors but a handful of articles have hundreds of editors is an example of *article centralization*. In the upper left of editor degree distributions (Figures 4a, 4d), we observe the top-ranked user edited nearly every article in a given year, whereas in the lower right we observe hundreds of users editing only a single article. Nonbreaking articles in different years exhibit remarkable similarity in editor centralization. The distributions for breaking articles in earlier years from 2001 to 2006 are truncated, which is largely an artifact of there being fewer breaking articles in these years, thus a smaller total population of editors. However, the editor centralization for breaking articles in later years has a similar intercept, shape, and slope as the nonbreaking articles' distributions. This is evidence that breaking articles do not differ significantly from nonbreaking articles' editor centralization.

Article degree distributions (Figures 4b, 4e) exhibit similar patterns where the top-ranked article receives contributions from almost all of the editors in a year whereas




Figure 4b. Article degree distributions, by year for breaking articles




Figure 4c. Edge weight distributions, by year for breaking articles




Figure 4d. Editor degree distributions, by year for nonbreaking articles




Figure 4e. Article degree distributions, by year for nonbreaking articles




Figure 4f. Edge weight distributions, by year for nonbreaking articles

the lowest-ranked articles receive contributions from only a few editors. Again, non-breaking articles exhibit remarkable similarity across the years, whereas breaking articles' distributions are less centralized in early years compared to later years. However, the centralization of breaking articles in later years exhibits similar intercepts, slopes, and shapes as the centralization of nonbreaking articles, which is also evidence that article centralization does not differ significantly between breaking and nonbreaking article types. Finally, the distribution of edge weights captures the *contribution centralization*, or how many times a single editor contributed to a single article (Figures 4c, 4f). These also follow long-tailed behavior; a handful of editors make hundreds of contributions to a single article, but the vast majority of editors make only a single contribution to an article when they edit. Nonbreaking articles exhibit consistency in the distribution across years, and again breaking articles in later years converge to the same distribution.

Figure 5 plots the article degree correlations of breaking and nonbreaking articles in each year. In bipartite networks, the article degree correlation reflects the extent to which articles having many contributors also have contributors who have revised many other articles. Both breaking and nonbreaking articles exhibit a positive correlation (also termed degree assortativity), as articles having many editors tend to be revised by prolific editors revising many other articles, whereas articles with fewer editors tend to have editors revising fewer other articles (Newman, 2002). The




Figure 5a. Article degree correlation, by year for breaking articles

degree correlations for nonbreaking articles are relatively stable for all articles since 2001, whereas there has been an intensification and stabilization of this assortative pattern on breaking articles since 2001 (as seen by the increased level of assortativity over time).

Taken together, the results from Figure 4 suggest mixed support for Hypothesis 1a, which predicted the concentration of editors of breaking articles would differ significantly from that of nonbreaking articles. There is little evidence that collaborations involving breaking news articles coordinate this work by centralizing or decentralizing their work to a greater extent than for nonbreaking articles. However, as we show in the third analysis below, this observation is an artifact of breaking news articles becoming nonbreaking articles in the long run; significant differences in edit activity centralization or distribution are present in the immediate aftermath of the event, but these differences become diluted as the collaboration reverts to less high-tempo work later on. There is also mixed support for Hypothesis 1b, which predicted the breaking articles in different years would exhibit similar concentrations of editor activity. Breaking articles have seen substantial growth in editorship, which has changed the structure of their collaborations over the years whereas nonbreaking articles remain relatively stable. However, the distribution of well-connected articles, editors, and




Figure 5b. Article degree correlation, by year for nonbreaking articles

links within these networks are remarkably similar across time, and both types of articles exhibit strong assortative degree mixing patterns.

Organizational Regeneration

To test the extent to which breaking news collaborations are instances of organizational regeneration, we measured whether editors active in a given year were also active in other years. Figures 6a and 6b illustrate the overlap among editors between years for both breaking and nonbreaking articles. As expected, much of the editor overlap for breaking articles occurs among editors in adjacent years (cells adjacent to the diagonal): Among editors who contributed to breaking articles about incidents in 2008, 27.2% also contributed to breaking articles about incidents in 2007. Conversely, among the editors who contributed to breaking articles about incidents in 2007, 26% also contributed to breaking articles about incidents in 2008. The vertical axis for an incident’s year can thus be interpreted as the *persistence* of that cohort across subsequent incident years, whereas the horizontal axis for an incident’s year can be interpreted as the *size* of the cohort relative to the other cohorts. Of the editors contributing to breaking articles about events in 2011, 1.4% persisted from making changes to

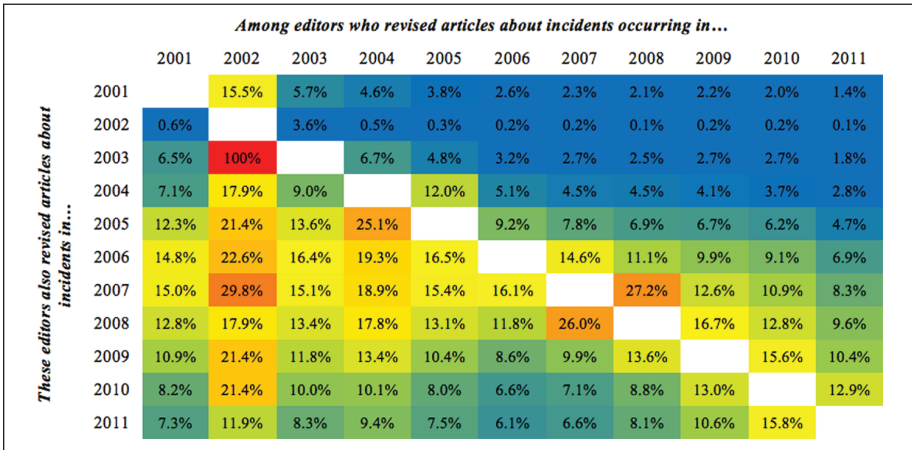


Figure 6a. Editor overlap for breaking articles (e.g., 15.5% of editors who contributed to 2002 breaking articles also contributed to 2001 breaking articles)

articles that had been breaking in 2001, whereas 7.3% of editors on articles that had been breaking in 2001 were still contributing in 2011.

Among breaking articles, there is a marked distinction in the overlap between editors who revised breaking articles about events before 2006 and editors who revised breaking articles about events since 2006. Historically, this distinction corresponds to the era of Wikipedia’s broader adoption, but it also reveals an interesting asymmetry in activity on these two types of articles. Editors of breaking articles about incidents before 2006 demonstrate substantial persistence and make up a nontrivial percentage of editors in subsequent years. This is largely attributable to the growth of the user population and concentration of activity on latter-day articles having hundreds or thousands of editors. However, Figure 6 suggests this asymmetry is notably absent among nonbreaking articles across all years. The persistence of editing cohorts across incident years as well as the stability of their presence on these articles suggest the existence of a core of editors revising nonbreaking articles about incidents across different years as well as the relative immunity of nonbreaking articles’ collaboration structure from growing the number of Wikipedia editors over time. These results reinforce the previous observation that nonbreaking article coauthorship is highly stable compared to breaking article coauthorship. These findings again provide mixed support for Hypothesis 2.

Coauthorship Network Evolution

The prior analyses compared features of centralization and regeneration on breaking and nonbreaking articles about events that occurred between 2001 and 2011 while

		<i>Among editors who revised articles about incidents occurring in...</i>										
		2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
<i>These editors also revised articles about incidents in...</i>	2001		26.1%	15.9%	16.4%	17.4%	15.7%	14.2%	10.3%	10.7%	10.1%	13.9%
	2002	30.5%		16.7%	17.2%	18.0%	17.0%	14.8%	10.8%	11.3%	10.5%	13.7%
	2003	25.3%	22.6%		49.5%	48.4%	44.1%	40.0%	37.9%	43.4%	41.8%	14.6%
	2004	24.0%	21.4%	45.3%		48.3%	43.9%	39.3%	37.7%	42.8%	41.4%	14.8%
	2005	25.5%	22.6%	44.7%	48.6%		45.5%	40.6%	38.1%	43.8%	42.1%	15.0%
	2006	25.4%	23.6%	45.0%	48.8%	50.3%		42.6%	39.1%	44.7%	42.8%	16.3%
	2007	25.0%	22.4%	44.4%	47.7%	48.9%	46.4%		40.7%	45.7%	43.4%	17.1%
	2008	16.7%	15.0%	38.8%	42.0%	42.3%	39.2%	37.5%		45.8%	43.7%	17.9%
	2009	14.7%	13.3%	37.5%	40.3%	40.9%	37.8%	35.4%	38.6%		57.5%	17.1%
	2010	14.3%	12.7%	37.1%	40.1%	40.5%	37.3%	34.7%	37.9%	59.2%		31.7%
	2011	7.6%	6.4%	5.0%	5.6%	5.6%	5.5%	5.3%	6.0%	6.8%	12.3%	

Figure 6b. Editor overlap for nonbreaking articles (e.g., 26.1% of editors who contributed to 2002 nonbreaking articles also contributed to 2001 nonbreaking articles)

Wikipedia existed. In this section we add a third genre called historical articles about events that occurred between 1990 and 2001. Furthermore, the previous approaches looked at the aggregate network structure rather than network structure at different points in time. To test Hypothesis 3, we use the snapshot approach we described above to examine the size of all three networks at different stages of their development and track the evolution of these networks. The snapshot networks in Figure 7a plot the number of editors editing each of three classes of articles at each of the 12 time thresholds. The number of editors active on breaking articles (in red) is relatively stable across time, whereas nonbreaking (blue) and historical articles (green) see substantial gains in the number of collaborators as articles reach about the age of one year. A similar pattern is reproduced for the degree of articles across time in Figure 7b: Breaking articles tend to have the same number of contributors across time, but these articles have significantly more editors than either nonbreaking or historical articles until approximately 1 year of age. Activity on breaking articles is initially much more distributed (more editors involved) than on other types of articles (few editors involved), but in the long run all types of articles converge on similar levels of editor activity, which helps explain the similarity of the overall degree and weight distributions observed in the first analysis.

Changes in editor and article clustering patterns across time and article types in Figures 8a and 8b are illustrative of patterns of regeneration as they point to significant differences in the extent to which articles share editors in common and editors share articles in common. Clustering among editors in Figure 8a is initially significantly higher among breaking articles than other article types, suggesting that the early editors of breaking articles have a tendency to collaborate together on many of the same breaking articles. In contrast, the early editors of nonbreaking and historical articles

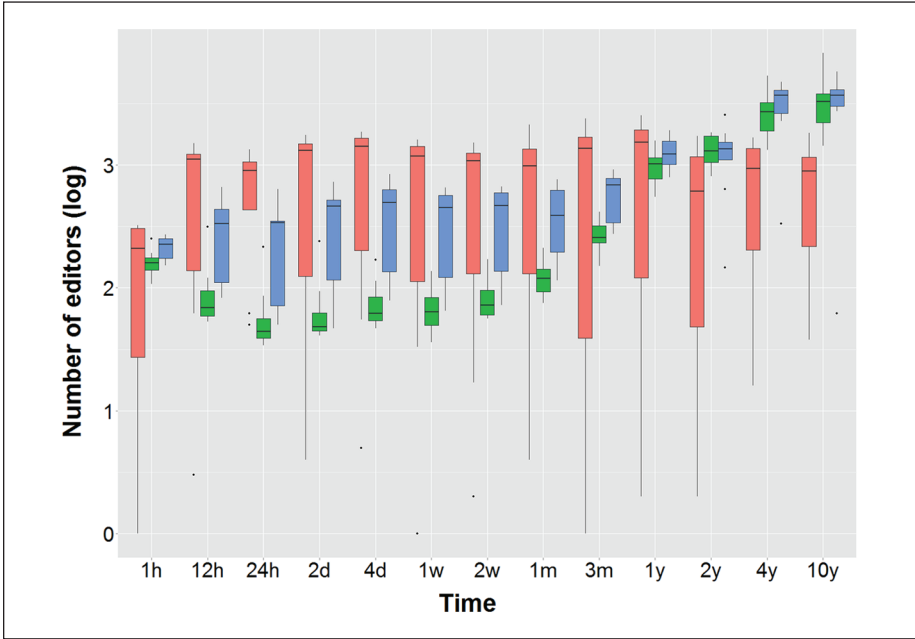


Figure 7a. Number of editors in the snapshot networks for each time threshold for breaking articles (red), nonbreaking articles (blue), and historical articles (green)




Figure 7b. Average article degree in the snapshot networks for each time threshold




Figure 8a. Editor clustering in the snapshot networks for each time threshold for breaking articles (red), nonbreaking articles (blue), and historical articles (green)




Figure 8b. Article clustering in the snapshot networks for each time threshold

are less clustered, but as time goes on these editors see dramatic gains in clustering and eventually exhibit similar levels of clustering as breaking articles after a year. Clustering among articles in Figure 8b is high in the first hour for all types of articles, suggesting each type of article has a tendency to initially share many editors in common. However, the article clustering drops off significantly for breaking articles, suggesting that editors making contributions between the ages of 12 hours and approximately 1 month are predominately new to editing articles of that type, but is relatively stable for the remainder of the article's history. Substantively, this means early editors of breaking news articles are unlikely to contribute at similar stages of other breaking news articles. In contrast, historical and nonbreaking articles see a more gradual attrition among articles sharing editors, suggesting that editors making changes to one article are likely to make changes to other similar articles. These different dynamics likely contribute to the observed differences in article overlaps: Organizational regeneration of prior collaborators is present at the outset of breaking news articles, but this feature is washed out over time as more editors contribute.

To test Hypothesis 3, we examined the extent to which breaking article collaborations are more cohesive at different stages of their development than either nonbreaking or historical articles. We measure the number of editors and articles in the largest connected component (LCC), which reflects the extent to which editors and articles are indirectly linked to each other. For networks with few editors or articles in the LCC at a given point in time, this suggests collaborations are more atomistic as editors revise articles independently of other editors who have a history of contributing to other articles in this domain. Alternatively, networks with many articles and editors in the LCC at a given point in time reflect coherent collaborations in which the editors work together on many of the same articles. The distributions of the editors in the LCC and articles in the LCC for each of the three types of articles across time in both the cumulative and snapshot networks are plotted in Figure 9a and 9b.

Unlike other article types, editing activity on breaking article collaborations (red) coheres into a large connected component within the first day of activity on these articles. The immediate emergence of this LCC for breaking articles requires the presence of editors who jointly revise many different breaking articles in the hours after these articles are created. Notably, this distribution includes breaking articles for all years between 2001 and 2011. This pattern contrasts with nonbreaking and historical articles (blue and green, respectively), in which activity is initially isolated but coalesces into a giant component after a year or more. This same pattern is also borne out looking at the snapshot networks where breaking news collaborations reliably have more activity in the largest component at every time threshold until the articles are approximately a year old. These results suggest breaking articles' collaboration structures significantly differ from those of both nonbreaking and historical articles throughout the first year, after which all article types exhibit similar tendencies for their activity to cohere into a single giant component. The convergence of all article types after a year suggests the collaboration dynamics of breaking articles are largely driven by the proximity to the event itself, but all types of articles cohere into large, connected collaborations after a year. These findings provide strong evidence for




Figure 9a. Articles in the largest connected component (LCC) for the cumulative network for breaking articles (red), nonbreaking articles (blue), and historical articles (green)




Figure 9b. Editors in the LCC for the cumulative network for each article type and threshold time

Hypothesis 3 that the dynamics of breaking articles are distinct from nonbreaking and historical Wikipedia articles at different stages of their life cycle.

Discussion

Our analysis offers new insights into how the self-organization of peer production systems differs from traditional patterns of collaboration in face of uncertainty and surprise. Comparing the centralization of editors and articles within breaking and nonbreaking Wikipedia article collaborations, both types of collaborations exhibit remarkable similarities in their density, clustering, and distribution of editor and article connectivity in the long run. We found evidence of repeat editorship across events in different years, but nonbreaking articles have a much greater tendency to rely on similar sets of editors, whereas repeat coauthors made up a substantially smaller portion of the population on breaking articles. Although these findings contradicted Hypotheses 1 and 2 to some extent, examining the time evolution of breaking and nonbreaking articles' collaborations revealed significant differences in their structures over time.

Wikipedia's breaking articles provide large-scale and longitudinal logs of user behavior, which allows us to analyze the structure and dynamics of high-tempo, online collaboration. Breaking news article collaborations operate under conditions of simultaneous uncertainty and time pressure that create tensions over whether to centralize or distribute activity among other group members. Whereas previous experimental studies of task performance have examined small groups completing arbitrary tasks, we examined knowledge collaborations "in the wild" with groups involving dozens or hundreds of individuals on substantially more complex and dynamic work. We find evidence that decentralization of activity is a prevailing tendency of online peer production groups during the most acute phases of these high-tempo collaborations, followed by a regression toward the interaction patterns typically found on nonbreaking and historical articles.

This study provides a theoretically motivated empirical basis for understanding the structure and dynamics of rapid online self-organization in sociotechnical systems. Our findings challenge assumptions in prevailing organizational theories about high-tempo collaboration that predominately examine physically colocated teams in which roles can be assigned and tasks coordinated by encoding specialization into material artifacts such as differentiated uniforms and routinized through shared professional norms. We find evidence that group members in high-tempo online collaborations differentially pattern their interactions during the most acute phases of article development as compared to nonbreaking and historical articles. The immediate emergence of connected components in the early stages of breaking article collaborations and the stability of breaking article editors' interactions with each other over time point to the presence of role specialization in editing breaking articles. Although our study did not qualitatively examine the practices adopted by editors of breaking news articles, our

findings suggest coordination on breaking articles proceeds from the regeneration of social roles, norms, and expectations among users who previously worked together (Bechky, 2006). Although the existence of these differentiated social roles validates prior work (Welser et al., 2011), the implicit or explicit processes by which these roles are negotiated and the social or cognitive channels through which they are shared remain unclear.

The size of the data in this corpus precluded any systematic examination of the situated practices editors employed to coordinate this work. The processes by which roles are negotiated and enacted, dependencies and expectations are negotiated and distributed, and practices are translated and adapted across different breaking article collaborations are rich domains for future inquiry. In particular, this class of breaking news articles presents a novel lens for future work to reevaluate extant conceptualizations of “communities of practice” as an integrating theory of socialization and coordination in online communities. Whereas communities of practice emphasize colocated and tightly knit groups practicing together long enough to develop mutual and shared understandings through sustained interactions and coordination (Bryant, Forte, & Bruckman, 2005), breaking article coauthorship is characterized by temporary and distributed work on immaterial artifacts among unfamiliar collectivities of loosely related individuals. A “collectivities” approach emphasizing how individuals’ knowledge is exchanged and competencies are integrated in high-tempo online collaborations may help reconcile how work proceeds when peripheral participation and deference to tenure are impracticable (Lindkvist, 2005).

The existence of diverse genres of Wikipedia articles about breaking news events, as well as collaborators who work across articles, provides a unique showcase to reflect on the possibilities of peer production and open collaboration in online communities. Ensuring the stability of the community of contributors and motivating sustained contributions over time is paramount to the success of many online communities, but participation in online communities does not always occur under conditions of stasis. Wikipedia’s breaking article collaborations rapidly accommodate and socialize large influxes of participants attempting to make sense of unexpected events by balancing competing interests to support openness, flexibility, and autonomy against institutional needs for structure, norms, and socialization over very different time scales (Keegan et al., 2011). More broadly, Wikipedians’ commitments to the synthesis and dissemination of timely, neutral, and reliable information about current events provide a case to reflect on how new forms of organizing characterized by mass collaboration and peer production are giving rise to new forms of participatory journalism.

Acknowledgments

The authors would like to thank members of the Collaborative Technology Laboratory and Science of Networks in Communities laboratory for their feedback and support, and Nick Bennett and Lynn Cherny for assistance and advice with data analysis and visualization scripts.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Funding was provided through National Science Foundation Grants 0838564, 0904356, and 1010904.

References

- Aral, S., & Alstynne, M. V. (2011). The diversity-bandwidth trade-off. *American Journal of Sociology, 117*(1), 90-171.
- Argote, L., Turner, M. E., & Fichman, M. (1989). To centralize or not to centralize: The effects of uncertainty and threat on group structure and performance. *Organizational Behavior and Human Decision Processes, 43*(1), 58-74.
- Bechky, B. A. (2006). Gaffers, gofers, and grips: Role-based coordination in temporary organizations. *Organization Science, 17*(1), 3-21.
- Birnholtz, J. P., Cohen, M. D., & Hoch, S. V. (2007). Organizational character: On the regeneration of Camp Poplar Grove. *Organization Science, 18*(2), 315-332.
- Brown, S., & Eisenhardt, K. (1997). The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations. *Administrative Science Quarterly, 42*(1), 1-34.
- Brown, T., & Miller, C. (2000). Communication networks in task-performing groups. *Small Group Research, 31*(2), 131-157.
- Bryant, S. L., Forte, A., & Bruckman, A. (2005). Becoming Wikipedian: Transformation of participation in a collaborative online encyclopedia. In *Proceedings of the 2005 ACM conference on supporting group work* (pp. 1-10). New York, NY: ACM.
- Burt, R. S. (2004). Structural holes and good ideas. *American Journal of Sociology, 110*(2), 349-399.
- Cohen, N. (2007). The latest on Virginia Tech, from Wikipedia. *New York Times*, April 23. Retrieved from <http://www.nytimes.com/2007/04/23/technology/23link.html>
- Deuze, M., Bruns, A., & Neuberger, C. (2007). Preparing for an age of participatory news. *Journalism Practice, 1*(3), 322-338.
- Faraj, S., & Xiao, Y. (2006). Coordination in fast-response organizations. *Management Science, 52*(8), 1155-1169.
- Ferron, M., & Massa, P. (2011). WikiRevolutions: Wikipedia as a lens for studying the real-time formation of collective memories of revolutions. *International Journal of Communication, 5*, 1313-1332.
- Jones, C., Hesterly, W., & Borgatti, S. (1997). A general theory of network governance: Exchange conditions and social mechanisms. *Academy of Management Review, 22*(4), 911-945.
- Keegan, B. (2011, May). *Breaking news, breaking planes, and breaking hearts: Psycholinguistic dynamics and sensemaking in collaborative accounts of catastrophe*. Paper presented at the International Communication Association meeting, Boston, MA.

- Keegan, B., Gergle, D., & Contractor, N. (2011). Hot off the wiki: Dynamics, practices, and structures in Wikipedia's coverage of the Tōhoku catastrophes. In *Proceedings of the 2011 symposium on wikis and open collaboration* (pp. 105-113). New York, NY: ACM.
- Keegan, B., Gergle, D., & Contractor, N. (2012). Do editors or articles drive collaboration?: Multilevel statistical network analysis of Wikipedia coauthorship. In *Proceedings of the 2012 ACM conference on computer supported cooperative work* (pp. 427-436). New York, NY: ACM.
- Kittur, A., Chi, E., Pendleton, B., Suh, B., & Mytkowicz, T. (2007). Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. In *Proceedings of the 2007 ACM conference on human factors in computing systems* (pp. 19-28). New York, NY: ACM.
- Kittur, A., & Kraut, R. (2008). Harnessing the wisdom of crowds in Wikipedia: Quality through coordination. In *Proceedings of the 2008 ACM conference on computer supported cooperative work* (pp. 37-46). New York, NY: ACM.
- Kittur, A., Lee, B., & Kraut, R. (2009). Coordination in collective intelligence: The role of team structure and task interdependence. In *Proceedings of the 2009 ACM conference on human factors in computing systems* (pp. 1495-1504). New York, NY: ACM.
- Latapy, M., Magnien, C., & Vecchio, N. D. (2008). Basic notions for the analysis of large two-mode networks. *Social Networks*, 30(1), 31-48.
- Lindkvist, L. (2005). Knowledge communities and knowledge collectivities: A typology of knowledge work in groups. *Journal of Management Studies*, 42(6), 1189-1210.
- Meyerson, D., Weick, K. E., & Kramer, R. M. (1996). Swift trust and temporary groups. In T. Tyler & R. Kramer (Eds.), *Trust in organizations: Frontiers of theory and research* (pp. 166-195). Thousand Oaks, CA: Sage.
- Mituzas, D. (2011). *Wikipedia article traffic statistics*. Retrieved from <http://stats.grok.se/>
- Newman, M. E. J. (2002). Assortative mixing in networks. *Physical Review Letters*, 89(20), 208701.
- Panciera, K., Halfaker, A., & Terveen, L. (2009). Wikipedians are born, not made: A study of power editors on Wikipedia. In *Proceedings of the 2009 ACM conference on supporting group work* (pp. 51-60). New York, NY: ACM.
- Pentzold, C. (2009). Fixing the floating gap: The online encyclopaedia Wikipedia as a global memory place. *Memory Studies*, 2(2), 255-272.
- Quarantelli, E., & Dynes, R. (1977). Response to social crisis and disaster. *Annual Review of Sociology*, 3, 23-49.
- Seward, Z. (2009). *Here's the AP document we've been writing about*. Retrieved from <http://www.niemanlab.org/2009/08/heres-the-ap-document-weve-been-writing-about/>
- Shklovski, I., Palen, L., & Sutton, J. (2008). Finding community through information and communication technology in disaster response. In *Proceedings of the 2008 ACM conference on computer supported cooperative work* (pp. 127-136). New York, NY: ACM.
- Starbird, K., Palen, L., Hughes, A., & Vieweg, S. (2010). Chatter on the red: What hazards threat reveals about the social life of microblogged information. In *Proceedings of the 2010 ACM conference on computer supported cooperative work* (pp. 241-250). New York, NY: ACM.
- Weick, K., & Roberts, K. (1993). Collective mind in organizations: Heedful interrelating on flight decks. *Administrative Science Quarterly*, 38, 357-381.

- Welser, H. T., Cosley, D., Kossinets, G., Lin, A., Dokshin, F., Gay, G., & Smith, M. (2011). Finding social roles in Wikipedia. In *Proceedings of the 2011 iConference* (pp. 122-129). New York, NY: ACM.
- Wikipedia. (2011, April). *Wikipedia statistics: English—Zeitgeist, 2010*. Retrieved from <http://stats.wikimedia.org/EN/TablesWikipediaEN.htm>
- Wilkinson, D. M. (2008). Strong regularities in online peer production. In *Proceedings of the 2008 ACM conference on electronic commerce* (pp. 302-309). New York, NY: ACM.

Bios

Brian Keegan is a post-doctoral research fellow in computational social science at Northeastern University. Keegan uses network analysis to understand the structure and dynamics of online collective action such as Wikipedia, political information seeking, and crowdsourced creativity.

Darren Gergle is an Associate Professor at Northwestern University where he directs the Collaborative Technology Laboratory (CollabLab). His research investigates group communication in mediated environments and the design of tools to support distributed collaboration.

Noshir Contractor directs the Science of Networks in Communities (SONIC) Research Group at Northwestern University. He is investigating factors that lead to the formation, maintenance, and dissolution of dynamically linked social and knowledge networks in communities.