

Do Editors or Articles Drive Collaboration? Multilevel Statistical Network Analysis of Wikipedia Coauthorship

Brian Keegan, Darren Gergle, Noshir Contractor

School of Communication, Northwestern University

2240 Campus Drive, Evanston, IL 60201

{bkeegan, dgergle, nosh}@northwestern.edu

ABSTRACT

Prior scholarship on Wikipedia's collaboration processes has examined the properties of either editors or articles, but not the interactions between both. We analyze the coauthorship network of Wikipedia articles about breaking news demanding intense coordination and compare the properties of these articles and the editors who contribute to them to articles about contemporary and historical events. Using p*/ERGM methods to test a multi-level, multi-theoretical model, we identify how editors' attributes and editing patterns interact with articles' attributes and authorship history. Editors' attributes like prior experience have a stronger influence on collaboration patterns, but article attributes also play significant roles. Finally, we discuss the implications our findings and methods have for understanding the socio-material duality of collective intelligence systems beyond Wikipedia.

Author Keywords

Wikipedia; collaboration; network analysis; coauthorship; exponential random graph model; ERGM; socio-material

ACM Classification Keywords

H.5.3 Theory and Models, Computer-Supported Cooperative Work

General Terms

Human Factors; Measurement; Theory

INTRODUCTION

"What makes Wikipedia work?" is a pervasive question in the literature on computer-supported cooperative work and social computing. The motivations of editors, design features of the community, affordances of peer production, and other latent social processes interact in complex ways at multiple levels to enable and sustain this massive collaboratively authored online encyclopedia.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW'12, February 11–15, 2012, Seattle, Washington, USA.

Copyright 2012 ACM 978-1-4503-1086-4/12/02...\$10.00.

Previous research makes clear that editors of Wikipedia articles fulfill distinct and diverse collaboration roles and different types of articles employ different forms of coordination [1, 2]. However, extant scholarship has not examined the interaction between these features: how do editors with particular skills self-organize around articles requiring different forms of collaboration?

To answer this question we examine a "boundary condition" for self-organization on Wikipedia. We compare the coauthorship of Wikipedia articles about current and breaking news events such as commercial airline disasters to topically similar articles about historical airline disasters. Articles about breaking news are coauthored under "high tempo" conditions which demand unique forms of coordination to manage interdependencies. We analyze the coauthorship networks of high and low tempo articles as well as the attributes of editors who contribute to them.

We review existing scholarship on the duality of Wikipedia as both user action embodied in artifacts and processes which support collaboration. Integrating this review, we develop a multi-theoretical, multi-level model describing how features of editors, articles, and interactions between both influence collaboration structure [3, 4]. Our findings suggest that while the features of articles and attributes of editors both influence structure, editors' experience more strongly governs the types of editors they collaborate with and the types of articles they work on. These findings have implications for how task demands intersect with user attributes to structure self-organizing collaborations.

We also make a methodological contribution by demonstrating how a class of statistical methods called p*/exponential random graph models (p*/ERGMs) enable multi-level network analysis. We specify statistical parameters which correspond to processes operating at each of the article and editor levels to disentangle which are more influential on collaboration structure. We discuss the implications p*/ERGM methods have for analyzing and comparing multi-level social interactions in other domains.

BACKGROUND

As is the case with many online communities, the majority of contributions to Wikipedia come from a fraction of the entire user base [5, 6]. Despite this disparity in effort, online communities like Wikipedia are able to escape traps

such as the tragedy of the commons and social loafing owing to members' uses and gratifications [7], diverse motivations [8], affordances of peer production [9], and design features of the community itself [10]. These approaches generally emphasize the agency of individuals to form and realize their own goals. In the context of Wikipedia, editors who are motivated to fulfill particular roles like copyediting or vandal fighting [1, 11], are socialized into sub-communities with like-minded collaborators [12], and recognize or revert the contributions of other editors [13].

Peer-production communities are oriented around producing and maintaining an artifact such as an operating system or encyclopedia article. Although these artifacts are the agglomeration of individual users' actions, artifacts like Wikipedia articles are embedded within socio-technical systems which imbue them with innate agency. This material agency enables them to operate outside of the control of any single person and emerges from the network of user and system interactions [14]. In the context of a Wikipedia article, content added or reverted by other editors, markup language interpreted into style formatting by the MediaWiki software, and page protections enforced by administrators are examples of articles "acting" on their "own" outside of any one user's direct control.

Coauthorship patterns on Wikipedia thus foreground the duality of persons and material artifacts [14-16]: collaborations occur around articles exhibiting particular features but articles also emerge from the contributions of editors with distinct traits. The material agency of articles such as topic, quality, age, or number of contributors influences the types of editors who are capable of making further contributions to an article. For example, a featured article about Barack Obama probably will not preserve contributions from newly-registered editors identifying with the Tea Party movement. However, the properties of editors also influence the types of articles they choose to edit and maintain. The human agency of these actors manifests in attributes such as varying expertise, editing experience, and roles. For example, a college freshman who uploads photos about soccer players will be unlikely to take up correcting formulae on general relativity. This suggests the features and attributes of *both* articles and editors influence the self-organization of collaboration on Wikipedia.

While scholars have articulated rationales for how and why collaboration and social action emerge from both internal human agency (what we refer to as editor-focused attributes) as well as external social foci (what we refer to as article-focused features), the interaction between these two approaches has not been studied. We review and identify themes from each of these literatures to motivate a statistical analysis which allows us to "decouple" this duality and model the individual influences as well as interactions between editor attributes, article features, structural pattern of editors' contributions to articles, and

structural patterns of an article's contributions from editors. We advance CSCW scholarship by using a statistical approach that allows us to simultaneously examine how the features of articles, attributes of editors, and these network structures *all* influence the organizations of Wikipedia collaborations. While processes which occur at a single level of analysis (such as 1 and 2 below) certainly play a role, we expect the *interactions* between each of (1) an editor's attributes and his or her history of editing other articles, and (2) an article's features and its history of revisions from other editors, will provide a more complete account of self-organization on Wikipedia.

1. Editor based attributes. For example, do experienced editors contribute to more articles than non-experienced editors?
2. Article based features. For example, do breaking news articles involve more editors than traditional articles?
3. Editor-focused interactions with article features. For example, is an experienced editor more likely to contribute to breaking news articles if they previously contributed to *other breaking news articles* than a non-experienced editor?
4. Article-focused interactions with editor attributes. For example, are breaking news articles more likely to attract contributions from experienced editors if *other experienced editors* have also contributed than non-breaking news articles?

In the following sections we examine the prior literature that has typically focused either on article-focused features or editor-focused attributes and aim to consolidate this work using a multi-theoretical, multi-level modeling approach that allows us to describe how both features of articles, editors and the interactions between them influence the structure of collaborations on breaking news events about commercial airline disasters.

Article-focused: Task Coordination and Social Foci

The demands of coordinating coauthorship on Wikipedia articles vary substantially with the age of the article and the number of contributors to it [2, 17]. An additional dimension is the contemporary salience of an article. Like other forms of social media [18], current and breaking news events uniquely motivate editors to contribute to and collaborate in Wikipedia [19]. However, co-authoring an article about breaking news events like commercial airline disasters involves complex, time-sensitive, and highly interdependent tasks. In this section we review how the features of breaking news articles not only influence patterns of coauthorship, but how these article-focused features *interact* with the attributes of the editors who contribute to these articles.

Although prior research suggests the compounding coordination costs of many editors engaged in interdependent work will inhibit the development of high quality Wikipedia articles [2, 17], articles about breaking news complicate this assumption. On one hand, Wikipedia

articles about breaking news events are often perceived to be exemplars of timeliness, breadth, and reliability in the immediate aftermath of an event like the Virginia Tech massacre [20]. On the other hand, the concentration of new editor activity is also densest while the article is less than 24 hours old and being intensively developed [19].

The popularity and quality of these articles in spite of these constraints poses a paradox in which “breaking articles” remain high quality in spite of the number of editors attempting to make simultaneous contributions with incomplete information and no centralized coordination. Examining the features of only articles or editors may present an incomplete picture and the novel coordination processes which enable these articles to be rapidly authored but also high quality likely emerge from the interactions between features of the article and attributes of its editors.

We argue these breaking news articles belong to a class of *high-tempo collaborations* characterized by non-routine and extremely urgent work, abrupt consequences, and intense attention. Coordination in these volatile environments demands high levels of heedful and interrelated action, knowledge integration, and information processing [21]. These “emergent response groups” are unique because group members have diverse motivations, mixed perspectives, varied resources to contribute, and substantial volition to come and go as they please. These factors contribute to unstable task definitions and the pursuit of potentially conflicting goals [22].

Members of these collaborations adapt by re-tailoring and sharing their particular expertise, emphasizing trust through action rather than credibility through expertise, and relying on narratives and knowledge artifacts to document actions taken [22]. What emerges from the on-going and repeated interactions between both editors and the article as they expand, update, copy edit, and fight vandals is not only the content of the article but also an artifact narrating prior actions and decisions.

Thus, an article feature such as being breaking or non-breaking is an important variable for modeling the collaboration patterns of editors. The salience of the topic and demands of coordinating interrelated tasks makes breaking news articles *foci* which actively bring people together and shapes their collective action much more than articles about historical events which do not demand high-tempo collaborations [23]. Therefore, we expect that articles about breaking news events may attract more editors than Wikipedia articles about non-breaking news events. Thus, we expect:

H1 Article attributes like salience will co-vary with number of editors. Breaking news articles will attract more editors than non-breaking articles.

There are latent tendencies for breaking news articles to receive many contributions (discussed above) or experienced editors to simply be prolific or engaged in

many articles (discussed in the next section). However, considering the features of the article absent the attributes of the editors who contribute to it is necessarily incomplete. For example, the “sink or swim” coordination demands of a breaking news article or attempts to limit “dysfunction from diversity” [24] may predispose editors to only want to collaborate with other editors who exhibit similar characteristics or qualifications as themselves on a breaking news article. If other experienced editors are contributing this may be social proof about the collaboration and may attract other experienced editors. This would manifest as an article-focused homophily in which the attributes of the article result in similar kinds of editors collaborating [25].

H2 Article attributes like salience will co-vary with the types of editors who collaborate with each other. Experienced editors will co-author with other experienced editors on breaking news articles.

Editor-focused: Social Roles and Identity

An alternative rationale why collaborations around breaking news articles exhibit different processes of self-organization revolves around the attributes of the editors who contribute to these articles rather than the features of the articles themselves. Editor-focused attributes like experience can potentially explain why some editors contribute to more articles than others, but also how these editor attributes interact with article features and lead users to fulfill distinct roles contributing to particular types of articles.

Roles in social media manifest as behavioral regularities, structural position, social action, or self-identification. These roles form complex ecologies which are defined in relation to other roles such as substantive experts, technical editors, counter-vandalism, and community building [26]. Other typologies have identified the placeholders, completers, housekeepers, and shapers who contribute, integrate, and shape content on Wikipedia [27].

In general, roles are resources that enable individuals to adapt to new contexts by creating new structures as well as imitating behaviors that were previously successful [28]. Because dedicated Wikipedia editors have distinct but stable behavioral patterns [6], editors can be classified into distinct roles based on the distribution of their activity [1]. The patterning of interactions among editors inhabiting particular social roles across different breaking news article collaborations can potentially explain how breaking news articles are co-authored in spite of steep coordination costs.

Editors may have a particular interest in contributing to topical areas such as airliner disasters (the topic explored in this paper). Experienced editors in this domain may have deep expertise about the appropriate vocabulary and style for describing an accident or knowledge about relevant citations [12]. Thus, these editors’ may be fulfilling “caretaker” roles in which they edit many articles while less experienced editors specialize on fewer articles.




Figure 1: Visualization of p^* /ERGM attribute interaction parameters that capture varying attributes of both the editors and the articles in Wikipedia. Dark-red circles are expert editors, medium-red circles are apprentice editors, and light-red circles are non-expert editors. Dark-blue squares are breaking articles, medium-blue squares are contemporaneous articles, and light-blue squares are historical articles.

H3

Editor attributes like experience will co-vary with the number of articles edited. Experienced editors contribute to more articles than other editors.

Again, considering the attributes of editors separately from the features of the articles they edit presents an incomplete account by divorcing the role editors play from the resources to which they contribute. Experienced editors may self-identify as “breaking news editors” who preferentially edit these articles out of novelty or immediate gratification of contributing to an in-demand information artifact. Experienced editors preferring to edit only breaking news articles would manifest as a pattern of editor-focused homophily in which the attributes of editors cause them to contribute to similar kinds of articles. Thus, we expect:

H4

Editor attributes like experience will co-vary with the types of articles they edit. Experienced editors will be more likely to contribute to similar types of articles than dissimilar types of articles.

Alternative Explanations for Collaboration Structure

In addition to the hypothesized explanations, we expect a variety of alternative processes could account for differences in collaboration structures between articles covering breaking, contemporary and historical events. We control for the influence of these factors by including them as parameters in the model in addition to our hypothesized variables of interest (breaking news article and editor experience).

Article features such as the severity of a catastrophe, proximity to developed countries, evaluated quality, and article length are also likely to influence collaboration patterns. For example, in our corpus of breaking news events centering on airline disasters we expect that the number of fatalities and survivors of an accident will be strongly correlated with the amount of attention an article receives from Wikipedia editors. Wikipedians also exhibit a “self-focus bias” in which geographic proximity influences

why articles appear in one language but not another [29]. We expect that accidents occurring within or near developed countries will receive more attention from editors than accidents occurring in developing countries. Finally, a number of studies have identified that the number of editors, length of an article, and article quality are all correlated [30, 31]. We use both the article quality and article length as controls on the number of editors who contributed to the article.

Editor attributes such as tenure in the community and whether or not the editor is registered are also likely to influence the likelihood of editors making contributions to articles of various types. Editors who started editing earlier in Wikipedia’s history likely have greater familiarity with best practices and may be more involved in editing many articles [12]. We also expect registered editors are both highly motivated and more committed to the community, both of which lead them to make more contributions than non-registered editors [10].

OUR APPROACH

Establishing which of the collaboration processes is most influential requires a model accounting for the simultaneous contribution of each of these processes. However, owing to the methodological limitations of common types of network analysis, studies often only examine one level of analysis which accounts for little of the overall variance in the network. Analyzing the effects of network parameters interacting at different levels of analysis allow us an integrated test of complimentary and competing theories of how network structure emerges [3, 4].

While descriptive approaches and regression analyses serve valuable purposes for answering particular research questions, these approaches are limited to analyzing the properties in a “snapshot” of a network. These kinds of analyses cannot explain the endogenous processes of how the network structures itself nor the simultaneous influence of exogenous actor-level attributes on the network structure.

Statistical models like p^*/ERGMs extend the logic of multivariate logistic regression to relational data: the presence or absence of a binary tie in a network is an outcome variable predicted by a vector of independent variables called parameters. These parameters correspond to theoretically-motivated structural characteristics(s) we believe are more or less likely to occur in the distribution. Visualizations of these parameters are provided in Figure 1.

Because the likelihood that a network tie is present or absent in a network is not independently and identically distributed (IID) from other network ties' likelihoods, a statistical model must account for these dependencies [32]. In a p^*/ERGM , parameters allow the models to reflect dependencies on both endogenous tendencies for ties to exist because of other local structures in the network (e.g., popular articles continue to attract more editor attention) as well as exogenous attribute covariates (e.g., experienced editors prefer to work with other experienced editors). These methods allow us to test hypotheses about the tendency for ties to form as a result of the properties of the sending node, receiving node, as well as the presence and properties of other local ties and nodes.

The model produces a set of parameter estimates whereby estimates of zero indicate the modeled effect does not alter the likelihood from random chance, a positive parameter suggests the effect increases the likelihood of a tie, and a negative parameter implies the effect decreases the likelihood of a tie. Significance is tested using a t -ratio and concluded to be significant when the absolute value of this ratio exceeds a critical t -value of 1.96. Details about the specification, estimation, and simulation of p^*/ERGMs are beyond the scope of this paper but can be found in [32-35].

DATA, VARIABLES, AND METHODS

A variety of Wikipedia article genres like natural disasters, sporting events, and political scandals exhibit high tempo features which should require unique forms of coordination. We examine articles about commercial airline disasters for this study because these incidents occur with sufficient regularity to generate a large sample but are also "normal accidents" involving complex technology with prompt and serious consequences which make reliably notable events warranting coverage in Wikipedia.

Articles for the corpus were drawn from the "List of accidents and incidents involving commercial aircraft."¹ Our sample excludes hijackings and other instances of terrorism (such as the four flights involved in the September 11 attacks) because these incidents represent an archetype of catastrophe with distinct attention, salience, and causal attributions compared to "typical" accidents attributable to crew error, mechanical failure, or weather conditions. The resulting corpus includes 249 articles about

commercial airline disasters which occurred since January 1990 and through December 2010.

Based on the list of articles identified above, we developed a tool to extract and stitch together XML revision histories for each article using Wikipedia's "Special:Export" function.² These data include revision-level data about the article name and ID, editor name and ID, timestamp, content, and comments. Registered users ($n=6,462$) identified by names and unregistered users ($n=7,830$) identified by IP addresses each have unique IDs. 14,292 unique user accounts made 58,500 revisions to this corpus between September 16, 2001 and May 24, 2011.

Each revision's article and editor ID were recorded as a duple representing an editor's modification of an article. Because p^*/ERGMs estimate the binary presence or absence of a link rather than the weight or strength of a link, repeated editor-article duples were discarded creating a binary edgelist of 23,903 unique editor-article interactions. The edgelist was imported to the *statnet* statistical network analysis package in R for analysis using the *ergm* library [36]. The size of the resulting network required us to analyze the data on high-memory computing instances on Amazon Web Services' Elastic Cloud Compute (AWS EC2) infrastructure.

Bipartite Network Modeling

We conceptualize the Wikipedia revision data as a network consisting of a set of actors and a set of relationships among these actors. While traditional network analysis emphasizes unipartite or one-mode data where the relationships are between a single type of actor (i.e., people-to-people), two unique sets of actors actually exist in Wikipedia: editors and articles.

Because it is nonsensical for an article to edit another article or a user to edit another user, we structure our interaction data as a *bipartite graph* in which the nodes can be partitioned into exactly two mutually exclusive sets of actors such that ties only exist between sets and no ties exist within a set [37]. Thus, a link exists between an editor node, E , and an article node, A , if E made a contribution to A , but neither $E-E$ nor $A-A$ links can exist. This bipartite structure is alternatively referred to as a "two-mode" or "affiliation" network. A bipartite network is a natural approach for modeling collaboration because it explicitly models the duality of persons and groups: a link between the social actors as one level of analysis and the groups to which they belong as another [15]. We employ p^*/ERGM parameters designed specifically for bipartite networks [34, 35].

Node Attribute Construction

A variety of article and editor variables were extracted from either the revision histories or article content to provide covariates for control, analysis, and modeling.

¹ http://en.wikipedia.org/wiki/List_of_accidents_and_incidents_involving_commercial_aircraft

² <http://en.wikipedia.org/wiki/Special:Export>

- **Article attributes** – The date of the incident and timestamp of the first edit to the article were recorded. Based upon the difference between these times and Wikipedia’s creation in January 2001, three possible types of articles exist: *breaking articles* about events which are written within 72 hours of the incident ($n=93$), *contemporaneous articles* about events which happened while Wikipedia existed but were written well after the incident ($n=50$), and *historical articles* about events which predated Wikipedia’s existence and are thus written well after the event ($n=106$).
- **Other article attributes** – The *word count* excluding markup syntax is recorded as a continuous variable. The modal *article quality* as evaluated by associated WikiProjects is coded as one of seven categorical attributes ranging from Stub to Featured Article-class.
- **Incident attributes** – The count of *fatalities* and *survivors* are recorded for each event as controls for the severity and thus salience of a particular incident. A dummy control variable for *OECD location* coded whether or not the crash occurred within or off the coast of one of the 34 developed member nations of the OECD.
- **Editor attributes** – The total number of revisions an editor made *within the corpus* over the 10 year span of time was recorded as an edit count. This count was binned into three categories to capture topical editing experience: 1 to 3 contributions (4 is the mean) was *unexperienced* ($n=12,148$), 4 to 42 contributions (mean plus one standard deviation) was an *apprentice* ($n=1,992$), and 43 or more contributions was *experienced* ($n=152$). Editors who had registered accounts ($n=6,462$) were dichotomized from non-registered users ($n=7,830$). Finally, editors were partitioned into three classes: *early contributors* who made their first contribution to the corpus before 2008 ($n=5,366$), *middle contributors* who made their first contribution to the corpus from 2008 to 2010 ($n=6,224$), and *late contributors* who made their first contribution to the corpus between 2010 and the present ($n=2,702$).

P*/ERGM CAPTURES MULTILEVEL INTERACTIONS

We specify one large p*/ERGM which includes single level main effects and structural tendencies (summarized on Table 1) as well as multi-level attribute interaction parameters (summarized on Table 2) to test our four hypotheses. This method reproduces similar “main effects” findings for the control variables as previous studies: the likelihood of editors revising an article increases with every additional fatality associated with the incident, incidents occurring in or near OECD nations, higher quality, longer word counts, early contributors, and registered users.

We also specify endogenous structural parameters which control for the latent tendency for links to be created by chance (*edges*), the network-level tendencies for articles or editors to become highly centralized (*article* and *editor degree distribution*), and the local-level tendencies for articles to accumulate multiple editors (*multi-editor article*

Structural tendencies	
<i>Edges</i>	-2.626 (-217.0)***
<i>Multi-editor article tendency</i>	0.01321 (2.89)***
<i>Multi-article editor tendency</i>	0.08042 (36.9)***
<i>Article degree distribution</i> ($\alpha=2.5$)	-4.476 (-16.8)***
<i>Editor degree distribution</i> ($\alpha=0.25$)	4.616 (68.6)***
Article features	
Controls	
<i>Fatalities</i>	1.22E-03 (20.9)***
<i>Survivors</i>	1.61E-04 (1.40)
<i>Location: OECD</i>	0.145 (24.9)***
<i>Quality – Start</i>	0.158 (21.9)***
<i>Quality – C</i>	0.371 (56.1)***
<i>Quality – B</i>	0.366 (93.4)***
<i>Quality – GA</i>	0.309 (11.2)***
<i>Quality – FA</i>	0.127 (1.38)
<i>Word Count</i>	6.51E-05 (8.00)***
Main effects (Hypothesis 1)	
<i>Temporal Type – Contemporaneous</i>	-0.549 (-33.1)***
<i>Temporal Type – Historical</i>	-0.519 (-45.3)***
Editor attributes	
Controls	
<i>Registered</i>	0.700 (86.3)***
<i>Cohort – Middle contributor</i>	-0.137 (-6.87)***
<i>Cohort – Late contributor</i>	-0.154 (-3.45)***
Main effects (Hypothesis 3)	
<i>Experience – Apprentice</i>	-4.14 (-525.2)***
<i>Experience – Experienced</i>	-2.00 (-66.3)***

Table 1: “Main effect” p*/ERGM estimates (t-statistics).
Estimates are net of parameters in Table 2. Cells are shaded green for positive & significant estimates and red for negative & significant estimates.
 $p = *** < 0.001 < ** < 0.01 < * < 0.05$

tendency) as well as editors to contribute to multiple articles (*multi-article editor tendency*). The estimates for these parameters are listed in Table 1 as “structural tendencies.”

The estimate for the negative *edges* parameter reflects the log-likelihood of a network tie appearing entirely by chance and serves as the “intercept term” reflecting the density of the network if no other effects were present. This can also be interpreted as the “cost” of creating a tie which other structural tendencies, factors, and interactions will need to overcome. The positive *multi-article* and *multi-editor* structural parameters respectively reflect the latent tendency for articles to accumulate editors and for editors to edit many articles. The negative *article degree distribution* reflects the tendency for articles to avoid long-tailed degree distributions while the positive *editor degree distribution* captures the tendency for editors to have a very skewed distribution. These findings suggest a tie is most likely to form between articles with few co-authors and editors who have also edited many other articles.

We find evidence for H1 that breaking news articles are more likely to attract editors than contemporaneous or historical articles. Both of these article types are less likely to have ties to other editors than breaking news articles.

<i>Article-focused interactions</i> (Hypothesis 2)			
	<i>Breaking</i>	<i>Contemp.</i>	<i>Historical</i>
<i>Only Non-Experienced (P1)</i>	-0.0109 (-2.37)*	-0.0101 (-15.5)***	-0.0086 (-1.90)
<i>Non-Experienced & Apprentice (P2)</i>	-0.0153 (-3.33)***	-0.00122 (-2.59)**	-0.01449 (-3.02)**
<i>Only Apprentices (P3)</i>	-0.00317 (-0.68)	-0.03674 (-16.9)***	0.00166 (0.39)
<i>Non-Experienced & Experienced (P4)</i>	-0.01588 (-3.46)***	-0.02601 (-17.5)***	-0.01590 (-20.7)***
<i>Apprentice & Experienced (P5)</i>	0.00613 (1.10)	0.02869 (28.02)***	0.00415 (0.64)
<i>Only Experienced (P6)</i>	-0.03552 (-19.9)***	0.00578 (0.85)	0.00090 (0.23)
<i>Editor-focused interactions</i> (Hypothesis 4)			
	<i>Experienced</i>	<i>Apprentice</i>	<i>Non-Exp.</i>
<i>Only Breaking (P7)</i>	0.00552 (1.63)	0.11540 (28.3)***	-6.027 (-135.9)***
<i>Breaking & Contemp. (P8)</i>	-0.06421 (-20.7)***	0.01308 (2.50)*	-5.977 (-52.7)***
<i>Only Contemp. (P9)</i>	0.04338 (3.98)***	0.23140 (13.0)***	-5.116 (-32.7)***
<i>Breaking & Historical (P10)</i>	-0.07691 (-25.7)***	-0.01577 (-4.22)***	-6.136 (-81.7)***
<i>Contemp. & Historical (P11)</i>	-0.05648 (-18.5)***	0.08355 (25.7)***	-5.517 (-54.8)***
<i>Only Historical (P12)</i>	0.00056 (0.12)	0.15120 (45.5)	-5.468 (-84.6)***

Table 2: Editor-article “interaction” p*/ERGM estimates (t-statistics) capturing the tendency for the row title to edit/be edited by the column title. Estimates are net of parameters in Table 1. Labels in parentheses next to attribute names are references for attribute interaction parameters.
p = *** < 0.001 < ** < 0.01 < * < 0.05

However, the p*/ERGM estimates testing H3 invert our expectation that experienced editors would be more likely to contribute to many articles than non-experienced editors. The negative estimates imply apprentice and experienced editors are much less likely than non-experienced editors to edit many articles. Experienced editors (within our corpus) make repeated contributions to a few articles (again only within our corpus) rather than a few contributions to many articles. This specialization points toward rejecting H3.

Modeling Attribute and Structural Interaction Effects

p*/ERGM methods stand out from traditional regression approaches in their ability to model the interactions between local editor-article authorship structure, editor attributes, and article features with parsimonious and statistically-valid parameters. We employ bipartite p*/ERGM parameters like those visualized in Figure 1 and use them to model the complex interactions between article features and editor attributes. These feature and attribute estimates are summarized in Table 2.

There are two broad classes of interactions reflecting the two possible explanations for editor-focused or article-focused processes to influence the collaboration structure. We use editor experience and article temporal type as the interacting attributes. “Article-focused interactions” capture the tendency for articles possessing a feature such as

“breaking news article” to be edited by actors sharing similar or dissimilar attributes like experienced/apprentice/non-experienced editor. “Editor-focused interactions” capture the tendency for editors possessing an attribute like “experienced” to contribute articles sharing similar or dissimilar features like breaking/contemporary/historical article. Figure 1 visualizes six structural signatures for article-focused interactions with editor attributes on the left and six structural signatures for editor-focused interactions with article features on the right.

Within “Article-focused interactions”, we observe a general tendency for significantly fewer interactions than would be expected by chance between non-experienced editors and apprentices (P2) and experienced editors (P4). Coauthorship among non-experienced editors (P1) is likewise rarer than random chance. Likewise, there appear to be strong disincentives for apprentice editors to work with each other (P3) on articles about contemporaneous incidents. Despite our expectation that experienced editors would fulfill crucial roles in high tempo collaborations around breaking news events by intensively collaborating together, after controlling for variables such as the severity of the event, experienced editors work together on high tempo collaborations significantly less often than we would expect by chance (P6). Nevertheless, the other findings support H2 that the coordination demands of an article influence the tendency of editors with similar or dissimilar levels of experience to work together.

Within “Editor-focused interactions”, editors of all levels of experience are unlikely to contribute to both breaking and historical articles (P10). The lack of shared coauthorship on breaking and historical articles suggests these are very distinct sub-genres with limited interaction between each groups’ editors. However, apprentice editors have a tendency to edit diverse combinations of articles above and beyond the latent tendencies for editors to edit many articles. This offsets the main effect for apprentice editors to be unlikely to contribute to articles in general as well as the lack of a bridging role by experienced editors.

Experienced editors contribute to contemporary articles (P9) at a rate much greater than chance and make contributions to different types of articles (P8, P10, P11) at rates much less (respectively) than expected by chance. Contrary to our expectations, the effects of experienced editors’ sustained contributions to only breaking (P7) or historical articles (P12) are weak and non-significant. Highly experienced editors are instead characterized by deep and sustained involvement in a few articles rather than stewardship of many articles. Instead, it is the apprentice editors who appear to play a crucial role not only contributing to many articles but also acting as crucial brokers providing bridges within breaking (P7), contemporary (P9), and historical articles (P12) as well as between these article types (P8, P11). Again, these findings




Figure 2: Degree distribution with mean values from 10,000 simulated networks based on estimated model (in yellow) compared against observed values (in blue).

support H4 that an editor’s level of experience will influence the tendency for them to edit similar or dissimilar types of articles.

The magnitude of the coefficients for these editor-focused main effects and interaction parameters are generally larger than the article-focused parameters. This suggests the attributes and structural interactions focused on editors play a stronger role in explaining the presence and absence of links between editors and articles than the features and structures focused on articles. Thus, editor attributes are more influential on the self-organization of high tempo Wikipedia collaborations than the features of articles.

Confirming Goodness-of-fit by Simulation

The previous steps analyzed local-level processes but are these features sufficient to explain global network properties? We assess the model’s goodness-of-fit by simulating other networks based on this model and use the resulting distribution of networks to compare the properties of these generated networks to the observed network [32]. Using the `ergm` package’s “gof” function, we simulate a sample of 10,000 networks based on the p^*/ERGM in Tables 1 and 2 and measure fit using the degree distribution. Figure 2 plots the observed values (in blue circles) and distribution of simulated values (in yellow boxes) for the combined degree distribution for both modes of the network. We observe a good-fitting model because the observed distributions are almost completely bounded by the distributions from simulated networks.

DISCUSSION

Adopting a socio-material approach which recognizes the agency of both articles and editors to influence the self-organization of collaboration requires analyzing the interactions between both articles and editors. We incorporated editor attributes and article features by modeling their interactions as a bipartite graph and using p^*/ERGM methods. We expanded on previous approaches

by not only analyzing processes occurring at different levels of analysis but also by modeling the interactions between these levels in addition to controlling for potentially confounding processes. p^*/ERGM statistical models allowed us to disambiguate between the article-focused and editor-focused interactions by specifying a model which simultaneously incorporates each of these potentially confounding processes to assess the relative contribution of each to the network structure. This approach revealed new insights regarding how Wikipedia editors and articles self-organize in relation to one another.

Prior scholarship has either developed editor-focused accounts examining how editor attributes (e.g., experience) influence collaboration patterns [1, 26] or article-focused accounts of why some articles features (e.g., task coordination demands) lead to higher quality or more contributions [2, 17]. Our analysis is the first to simultaneously look at both levels of analysis to better understand the relationship each has on the self-organization of collaborations involving extreme coordination demands and varied editor experience. Our results suggest that editor experience and the features of articles in their contribution history have a stronger influence on the self-organization of the collaboration than article features like coordination demands and the attributes of editors who contribute to these articles. Our approach provides a more complete account of the processes which influence the structure of collaborations on Wikipedia than looking at the structure of network of just editors or articles.

Applying this approach to the paradox of how breaking news articles exhibit high quality despite steep coordination costs and varied editor experience, we unpacked how the attributes of editors have greater influence over this self-organization. These findings validated our hypotheses that not only are the coordination demands of articles matched with the number of editors who contribute to them (H1), but that coordination demands of certain article types also lead editors to seek or avoid other types of editors *depending on the type of editor* (H2).

However, as measured by both effect size and valence, the features of an article and its interactions with editor attributes play a secondary role in structuring the collaboration as compared to the attributes of editors and their interactions with article features. Although experienced editors exhibit a tendency toward concentrating their work in a few articles (H3), we found evidence that an editor’s level of experience leads them to also work on or avoid certain articles *depending on the type of article* (H4).

While previous work examined how editors’ varying level of expertise influenced how Wikipedia tools were used or other users were perceived and rewarded [12, 13], we demonstrate that editors’ patterns of contributions are mediated through their own intrinsic attributes, the coordination demands of an article, the kinds of articles they have contributed to in the past, and the types of editors

who also contributed to those articles. Unwinding these intricate dependencies is crucial for understanding the processes which contribute to the formation, maintenance, dissolution, and re-emergence of social and technological interactions in Wikipedia and other online communities and distributed organizations.

Implications

Wikipedia's coverage of breaking news events suggests that peer production systems are capable of operating far from conditions of stable task demands and community membership [19]. Our findings suggest that tasks which demand high tempo knowledge collaboration may benefit more from matching users to tasks based on their own experience level and history of contributions to similar tasks in the past rather than assembling a team solely by optimizing on the demands of the task or the experience of other members of a team. Recruiting members with diverse backgrounds and interests may increase group productivity [24], but our results suggest that special care should be paid to the particular configurations and combinations of interests rather than dimensionless indices of diversity.

Statistical models allow for a more parsimonious and theoretically-coupled representation of dense and complex network structures by capturing the local-level interaction tendencies as well as the emergent macro-level structure [35]. Moreover, complex dependencies in networks are difficult to make sense of with descriptive statistics, visualizations may not provide statistically valid inferences, and due to the differences in levels of analysis parameterizing and controlling for these complex dependencies is extremely difficult, if not impossible, with traditional OLS and even hierarchical regression (e.g., mixed model) techniques.

Like regression or other statistical approaches, p*/ERGM methods require specifying models with theoretically well-motivated parameters in addition to translating extant theoretical constructs into appropriate network parameters. Absent a theoretical rationale for model specification and appropriate controls, both types of models can recover spurious relationships. However, these approaches suggest scholars can pose more meaningful research questions about multi-level and multi-theoretical processes about self-organization in collective intelligence systems.

p*/ERGMs also allow comparative network analysis by examining the similarities of the processes which structure networks of very different size, scale, and context [38]. This analysis only looked at one particular sub-genre of articles about airline crashes, but it would be possible to estimate models for other topics with a breaking news component such as earthquakes, hurricanes, or sporting events. p*/ERGMs for each of these could be estimated and a meta-analysis performed to compare the collaboration practices across topics or even other collaboration systems.

Limitations and Future Work

p*/ERGM methods are computationally intensive and become even more so as both the complexity of the model and the number of nodes in the network increases. Although new "peta-scale" computational infrastructures may address these bottlenecks, for the time being extending p*/ERGM methods to very large networks containing tens of millions of nodes like the entire Twitter, Facebook, or Wikipedia graphs is impractical. However, well-motivated boundary specification and comparative analysis or sampling approaches combined with meta-analyses can make large-scale analysis more tractable [39].

The p*/ERGM we employed assumed the data was cross-sectional and thus omitted potential temporal dependencies such as a tendency for an editor to contribute *after* another editor contributes. Longitudinal models of network change and dynamics can also be specified [40]. Although, editors' social roles play a role in coordinating work, it may also be the case that articles can fulfill "roles" socializing editors into particular collaboration norms or introducing them to effective coordination practices. In light of the influence of these editor-focused attributes, future work should unpack how an editor's temporal "trajectory" of contributions influences the types of roles they fulfill across articles.

We encourage other researchers to adopt p*/ERGM methods to ask better questions about multi-level and multi-theoretical processes which influence communication patterns, knowledge sharing, and distributed collaboration in collective intelligence and other socio-technical systems.

ACKNOWLEDGMENTS

We would like to thank members of the CollabLab and SONIC groups for development assistance and feedback.

REFERENCES

1. Welser, H. T., Cosley, D., Kossinets, G., Lin, A., Dokshin, F., Gay, G. and Smith, M. Finding social roles in Wikipedia. In *Proc. iConference'11*, ACM (2011): 122-129.
2. Kittur, A., Lee, B. and Kraut, R. Coordination in collective intelligence: the role of team structure and task interdependence. In *Proc. CHI'09*, ACM (2009): 1495-1504.
3. Contractor, N. S., Wasserman, S. and Faust, K. Testing multitheoretical, multilevel hypotheses about organizational networks. *Academy of Management Rev.*, 31, 3 (2006), 681-703.
4. Monge, P. R. and Contractor, N. S. *Theories of Communication Networks*. Oxford U. Press, 2003.
5. Kittur, A., Chi, E., Pendleton, B. A., Suh, B. and Mytkowicz, T. Power of the Few vs. Wisdom of the Crowd: Wikipedia and the Rise of the Bourgeoisie. In *Proc. CHI'07*, ACM (2007).
6. Panciera, K., Halfaker, A. and Terveen, L. Wikipedians are born, not made: a study of power editors on Wikipedia. In *Proc. GROUP'09*, ACM (2009): 51-60.

7. Lampe, C., Wash, R., Velasquez, A. and Ozkaya, E. Motivations to participate in online communities. In *Proc. CHI'10*, ACM (2010): 1927-1936.
8. Rafaeli, S. and Ariel, Y. "Online motivational factors: Incentives for participation and contribution in Wikipedia." *Psychological Aspects of Cyberspace*. Cambridge U. Press, 2008.
9. Benkler, Y. *The wealth of networks: How social production transforms markets and freedom*. Yale U. Press, 2006.
10. Ren, Y., Kraut, R. and Kiesler, S. Applying common identity and bond theory to design of online communities. *Organization Studies*, 28, 3 (2007).
11. Geiger, R. S. and Ribes, D. The work of sustaining order in Wikipedia: the banning of a vandal. In *Proc. CSCW'10*, ACM (2010): 117-126.
12. Bryant, S. L., Forte, A. and Bruckman, A. Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. In *Proc. GROUP'05*, ACM (2005): 1-10.
13. Kriplean, T., Beschastnikh, I. and McDonald, D. W. Articulations of WikiWork: Uncovering Valued Work in Wikipedia through Barnstars. In *Proc. CSCW'08*, ACM (2008): 47-56.
14. Contractor, N., Monge, P. and Leonardi, P. M. Multidimensional Networks and the Dynamics of Sociomateriality: Bringing Technology Inside the Network. *Int'l Journal of Communication*, 5, (2011).
15. Breiger, R. L. The duality of persons and groups. *Social Forces*, 53, 2 (1974), 181-190.
16. Kane, G. C. and Alavi, M. Casting the net: A multimodal network perspective on user-system interactions. *Info. Sys. Research*, 19, 3 (2008), 253-272.
17. Kittur, A. and Kraut, R. E. Harnessing the wisdom of crowds in wikipedia: quality through coordination. In *Proc. CSCW'08*, ACM (2008): 37-46.
18. Palen, L. and Vieweg, S. The emergence of online widescale interaction in unexpected events: assistance, alliance & retreat. In *Proc. CSCW'08*, ACM (2008): 117-126.
19. Keegan, B., Gergle, D. and Contractor, N. Hot off the Wiki: Dynamics, Practices, and Structures in Wikipedia's Coverage of the Tohoku Catastrophes. In *Proc. WikiSym'11*, ACM (2011): 105-113.
20. Cohen, N. The Latest on Virginia Tech, From Wikipedia. *The New York Times* (2007).
21. Brown, S. and Eisenhardt, K. The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations. *Administrative Science Quarterly*, 42, 1 (1997), 1-34.
22. Majchrzak, A., Jarvenpaa, S. and Hollingshead, A. Coordinating expertise among emergent groups responding to disasters. *Org. Science*, 18, 1 (2007).
23. Feld, S. L. The focused organization of social ties. *American J. of Sociology*, 86, 5 (1981), 1015-1035.
24. Chen, J., Ren, Y. and Riedl, J. The effects of diversity on group productivity and member withdrawal in online volunteer groups. In *Proc. CHI'10*, ACM (2010): 821-830.
25. McPherson, M., Smith-Lovin, L. and Cook, J. M. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(2001), 415-444.
26. Gleave, E., Welser, H. T., Lento, T. M. and Smith, M. A. A conceptual and operational definition of 'social role' in online community. In *Proc. HICSS'09*, IEEE (2009): 1-11.
27. Yates, D., Wagner, C. and Majchrzak, A. Factors affecting shapers of organizational wikis. *J. of ASIST*, 61, 3 (2010), 543-554.
28. Baker, W. E. and Faulkner, R. R. Role as resource in the Hollywood film industry. *American J. of Sociology*, 97, 2 (1991), 279-309.
29. Hecht, B. and Gergle, D. On the localness of user-generated content. In *Proc. CSCW'10*, ACM (2010): 229-232.
30. Blumenstock, J. E. Size matters: word count as a measure of quality on Wikipedia. In *Proc. WWW'08*, ACM (2008): 1095-1096.
31. Hu, M., Lim, E. P., Sun, A., Lauw, H. W. and Vuong, B. Q. Measuring article quality in Wikipedia. In *Proc. CIKM'07*, ACM (2007): 243-252.
32. Robins, G., Pattison, P., Kalish, Y. and Lusher, D. An introduction to exponential random graph (p*) models for social networks. *Social Networks*, 29, 2 (2007), 173-191.
33. Robins, G., Snijders, T., Wang, P., Handcock, M. and Pattison, P. Recent developments in exponential random graph (p*) models for social networks. *Social Networks*, 29, 2 (2007), 192-215.
34. Faust, K. Centrality in affiliation networks. *Social Networks*, 19, 2 (1997), 157-191.
35. Wang, P., Sharpe, K., Robins, G. L. and Pattison, P. E. Exponential random graph (p*) models for affiliation networks. *Social Networks*, 31, 1 (2009), 12-25.
36. Handcock, M., Hunter, D., Butts, C., Goodreau, S. M. and Morris, M. statnet: Software Tools for the Representation, Visualization, Analysis and Simulation of Network Data. *J. of Stat. Software*, 24, 1 (2008).
37. Borgatti, S. P. and Everett, M. G. Network analysis of 2-mode data. *Social Networks*, 19, 3 (1997), 243-269.
38. Faust, K. and Skvoretz, J. Comparing networks across space and time, size and species. *Sociological Methodology*, 32, 1 (2002), 267-299.
39. Handcock, M. S. and Gile, K. J. Modeling social networks from sampled data. *Annals of Applied Statistics*, 4, 1 (2010), 5-25.
40. Snijders, T. A. B., Van de Bunt, G. G. and Steglich, C. E. G. Introduction to stochastic actor-based models for network dynamics. *Social Networks*, 32, 1 (2010), 44-60.