

Towards Quality Discourse in Online News Comments

Nicholas Diakopoulos and Mor Naaman
School of Communication and Information
Rutgers University
{diakop, mor}@rutgers.edu

ABSTRACT

With the growth in sociality and interaction around online news media, news sites are increasingly becoming places for communities to discuss and address common issues spurred by news articles. The quality of online news comments is of importance to news organizations that want to provide a valuable exchange of community ideas and maintain credibility within the community. In this work we examine the complex interplay between the needs and desires of news commenters with the functioning of different journalistic approaches toward managing comment quality. Drawing primarily on newsroom interviews and reader surveys, we characterize the comment discourse of SacBee.com, discuss the relationship of comment quality to both the consumption and production of news information, and provide a description of both readers' and writers' motivations for usage of news comments. We also examine newsroom strategies for dealing with comment quality as well as explore tensions and opportunities for value-sensitive innovation within such online communities.

Author Keywords

Commenting systems, online journalism, online discourse

ACM Classification Keywords

H.5.3 Group and Organization Interfaces: Asynchronous interaction; H.4.3 Communications Applications

General Terms

Human Factors

INTRODUCTION

Consumption of online news is becoming increasingly social and interactive, and now commonly includes facilities for sharing and commenting directly on news articles or via social network channels like Facebook or Twitter. A recent Pew survey found that 37% of online news users (and 51% of 18-29 year olds) think that commenting on news stories is an important feature to have, and that 25% of users have contributed their comments to online news articles [28]. Indeed, popular news and opinion site HuffingtonPost.com receives approximately three million comments every month. Such "ad-hoc" online communities are immensely

popular and pose ultra-large-scale interaction challenges associated with different interaction motives and with distilling expertise and knowledge from thousands of participants who may or may not know each other.

With the growing interest in the interactivity of news, the *quality* of the discussion anchored around online news stories is of paramount importance to news organizations wishing to stimulate public criticism, debate, and discussion while maintaining a credible community profile [16, 21]. While the largest online newspapers have the resources to police and effectively moderate their online space, smaller news organizations and local community newspapers are struggling with the quality and volume of the discourse that arises in their online communities [6].

In this paper we examine the complex interplay between the needs and desires of newspaper commenting system users with the functioning of different approaches toward managing quality within a journalistic context. Drawing primarily on newsroom interviews and user surveys, we characterize the comment discourse on SacBee.com, the online counterpart to the Sacramento Bee newspaper. We first identify the most salient quality issues on the SacBee.com site, and discuss the interplay and relationship of quality of discourse factors to both the consumption and production of news information. We then explicate an understanding of users' motivations for both reading and writing news comments. Finally, we examine newsroom strategies for dealing with quality in the comments and examine tensions and opportunities for innovation sensitive to the values of the newsroom culture.

RELATED WORK

Early work in Computer-Mediated Communication (CMC) has documented how a lack of status cues and social context can introduce unwelcome low-quality contributions into online communication systems. *Quality* in this context refers to a degree of excellence in communicating knowledge or intelligence and normatively includes notions of accuracy, reliability, validity, currency, relevancy, comprehensiveness, and clarity [30, 31]. In the realm of online comments low-quality contributions might include "flaming" and more impulsive remarks [14] and are often implicated with anonymity, with less anonymity linked to higher quality comments [15].

A variety of moderation systems and strategies have been studied as mechanisms for dealing with low quality

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 2011, March 19–23, 2011, Hangzhou, China.

Copyright 2011 ACM 978-1-4503-0556-3/11/03...\$10.00.

discourse in online commenting systems [4, 11, 17, 18, 34]. Best practices in online community building espouse posting clear community guidelines and expectations, and having dedicated moderators check and approve comments before the comments are made public (pre-moderation) [27]. In practice, moderation systems typically rely on either an explicitly defined moderator (pre- or post-moderation) or on aggregated crowd-based ratings (post-moderation). Some conceptions of moderation entail interjecting in a discussion in order to calm it or keep it on track [11]. In this paper we refer to such activity as *engagement*, with *moderation* left to mean the process by which abusive comments are evaluated and possibly removed from the system.

Crowd-based moderation has been shown to be an effective mechanism for filtering and ranking comment quality. Such systems work by aggregating community members' moderation judgments in order to indicate the quality of comments [18, 23]. Studies of how users personalize the results of such crowd-based moderation systems, such as that used on Slashdot, have also highlighted the heterogeneous goals of comment readers [17]. Our work extends the understanding of these heterogeneous goals through a qualitative description of users' goals and motives in both *reading* and in *writing* comments.

The process of moderation is imbued with additional meaning in the context of a journalistic news enterprise. Interviews with editors at a range of news institutions in the U.K. has shown moderation to be a way that journalists exercise their professional editorial roles as gate-keepers [13]. Ethnographies of newsrooms have suggested some of the contextual factors which moderate the adoption of more participatory features of new website, such as organizational hierarchies and professional newsroom culture [3, 25]. We contribute a deeper understanding of these factors in relation to online commenting systems and go further by incorporating the perspectives of comment readers and writers in order to derive value-sensitive design suggestions.

Despite the quality concerns and cultural barriers to implementation there is solid evidence of the journalistic (and economic) value of comments, including adding perspectives, insights, and personal experiences that can enrich a news story as well as enabling the tracking of user interests or getting sources and tips for future stories [6, 19, 21]. And while individual studies have looked qualitatively at editors' reactions to commenting systems [3, 6, 13] or quantitatively at user's comment reading and posting behavior [1, 17, 33], in this paper we synthesize across stakeholder perspectives to include editors, reporters, and moderators—as well as readers and commenters—in a study of factors and strategies that may be used to address comment quality in online news commenting systems.

STUDY OF NEWS COMMENTING

We undertook a study of the commenting system on the SacBee.com website with the goals of better understanding (1) the character and magnitude of the comment quality

issues as well as how they impact both users' and journalists' uses of the system, and (2) the functioning of different approaches toward managing quality within a journalistic context such as flagging, moderation, and reporter engagement. It was important for us to address these inquiries from a number of stakeholder perspectives including both user and newsroom views. To get the newsroom perspective we interviewed a diverse set of journalists in different roles including *editors*, *reporters*, and *moderators*. To get the users' perspective we deployed a survey on the website. Together with data gathered from the commenting system itself we were able to gain a holistic view of the system as well as to triangulate findings.

Data and Methods

For the interview portion of our study we conducted in-depth interviews with 18 people within the newsroom including editors, reporters, and moderators. All interviews were conducted face-to-face using a semi-structured protocol that was developed based on early conversations with members of the newsroom. Interviews were audio recorded and then transcribed. Sampling was purposive in order to include a variety of newsroom stakeholders and roles. Out of 18 interviewees, 8 held editor positions, 7 had reporter roles, and 6 served as moderators (for a total greater than 18 as 3 interviewees had both editor and moderator roles); 11 interviewees were male and 7 were female. Interviews ranged in length from 20 to 53 minutes ($M=32.5$, $SD=9.4$). The interview protocol included questions regarding internal uses and behavior around comments, failures and successes of the current system, and reactions to engagement and moderation strategies.

A survey of users was conducted between August 12 and August 24, 2009. The survey included both close- and open-ended items concerning behavior around reading, writing, and flagging comments. Questions also touched on issues of journalist engagement, moderation efficacy, and the perception of quality in the comments. Although quality is a multi-faceted and complex concept, early conversations in the newsroom led us to simplify the measurement of quality in the survey as the degree of perception of attacking, abusive, or otherwise offensive content.

Responses to the survey were solicited using a link posted on the homepage of SacBee.com, and via a blog post on the site. Of 792 people that loaded the page to take the survey, 390 people completed and returned the survey (49.2%). Respondents to the survey were predominantly male (61%) and well educated – 73.6% had undergraduate or graduate degrees. Most respondents were not subscribers to the print edition of the paper (65.2%). The median age was 45 years old ($SD = 11.34$). There may be a non-response bias to the survey results since the people who chose to click the link for the survey and who chose to submit survey responses may be those most interested in “fixing” the commenting system. Also these people may over-represent more engaged users since less engaged users of the site or commenting system may not have taken the time to respond.

Both the transcribed newsroom interviews and the open-ended survey responses concerning reading, writing, and flagging behavior were separately analyzed using a grounded-theory inspired methodology involving iterative coding of concepts and themes.

Finally, all of the 54,540 comments made on the site during the month of August 2009 were extracted from the moderation system. The data included comments that were deleted via moderation and comments made by blocked users. However, the data did not include any indication of comments that had been flagged but not deleted.

SYSTEM DESCRIPTION

We first characterize and provide some descriptive statistics of the commenting system used at SacBee.com in 2009, which was the focus of this study.

Functionality

Comments on SacBee.com are listed below each story in blocks of ten comments; users can click through numbered pages for each subsequent block of ten. Comments are sorted in reverse chronological order, but can also be sorted by chronological order, or recommendation score (high to low or vice versa). In order to leave a comment users are required to create an account on the site. The account registration procedure requires choosing a unique screen name and providing a valid email address, but does not require any other identifying information. Users that are uninterested in comments can click a box below any story to hide comments on all stories. On each comment the user can choose to “recommend” that comment, incrementing its recommend score, or they can “report abuse” on that comment. Comments are not threaded, so to respond to other comments users will typically reference the user name of the person they are responding to.

Like many online communities SacBee.com explicitly states its commenting policy. The policy includes warnings against posting profanities, hate speech, and personal attacks, as well as off-topic, repeat, or spam comments. It states that failure to abide by these rules can result in the deletion of comments or the banning of user accounts.

SacBee.com uses a post-moderation strategy to deal with abusive comments. If a user flags a comment as abusive this sends the comment to a queue that is accessible to editors and moderators of the site. Accompanying the flag there can also be a tag such as personal attack, advertising / spam, hate speech, vulgarity / obscenity, or copyright. If the comment receives 3 or more flags it is immediately hidden, whereas if it receives 1 or 2 flags it remains visible on the site until reviewed by a moderator. If a comment is deemed to be abusive by the moderators then it is removed from the site and only visible within the moderation system marked as *deleted*. Users who repeatedly abuse the system can be *blocked* by moderators; their comments remain visible to themselves but are hidden from other users.

Users and Usage

The number of users that commented on news articles is relatively small in comparison to the number of story readers. During August of 2009 there were 8,246 users who contributed at least one comment to the total 54,540 comments in our dataset. We observed a skewed distribution in commenting frequency: 46.5% of contributors wrote a single comment in the month, and 86.7% made less than 10 comments. There was a group of 162 heavy contributors that made more than 50 comments in the month and one user made 607 comments.

Of the 54,540 comments collected from the moderation system, 2,124 (3.9%) were marked as *deleted*. An additional 7,225 (13.3%) of comments were marked as *hidden* (i.e. they were submitted from *blocked* users). In other words, 45,191 (82.8%) of the comments were visible on the site. Of the 8,246 users that commented in August, 262 (3.17%) were blocked. Only 1.74% of users who commented less than 10 times were blocked, whereas 13.94% of users who commented more than 10 times were blocked ($\chi^2=354.45$, $p < .0001$). This suggests that frequency of commenting is a valuable indicator of what moderators would consider low quality discourse.

FINDINGS

We discuss the findings from our user survey and newsroom interviews below. We first present observations on the quality of discourse on the site and the challenges presented when it fails. Next, we discuss the perception of quality by readers and how that correlates with frequency of reading. We then examine the clash between the comment practices and journalism values. Moving from quality issues, we examine users' goals and needs in using commenting systems, and newsroom strategies in handling them.

Discourse Discontents

Despite efforts to keep the online discussion constructive and cordial, many respondents to our survey rated the comment discourse offensive in some way. The average agreement with the statement, “*I find some of the comments on SacBee.com to be offensive*” was 5.19 out of 7. One respondent to the survey summarized it thus, “*I read them because I'm curious what my fellow Californians think. I often don't read them because they can be downright ugly - racist, classist, rude and often ignorant.*”

Some journalists noted the impact of topicality on the nature of the discourse they observed in the comments. One editor put it thus, “*If you write about immigration here it just uncorks all sorts of comments, positions, opinions, vitriol*” (E4). Another editor recognized both the positive and negative influence topics can have on the discourse:

“If you have a warm fuzzy story, you don't have problems with those kinds of stories. If you have a policy debate kind of story, you generally get stuff within the confines of that. There are certain hot button subjects...it might be around social issues...that people feel very strongly about.” (E5)

But as another reporter explained, there is oftentimes a voice of reason in the comments, which can counteract more abusive responses:

“None of them fail completely. Usually there’s a voice of reason that interjects at some point and says you didn’t read the story well enough or you have to understand that this particular institution is facing XYZ or this person is not a bad person. Usually there’s a voice of balance in the posted comments...” (R5)

While in some cases the comments are functioning adequately, 5 of 7 reporters and 7 of 8 editors still voiced some fundamental concerns about comment quality. From our analysis the key themes in these concerns included issues of distortion, anonymity, personal attacks on sources or reporters, flaming, propagation of misinformation, and tarnishing the reputation of the paper. We elaborate on the journalistic implications of personal attacks and anonymity in the following section.

The Impact of Discourse Character on Journalism

The journalists we interviewed, as well as users that responded to our survey had objections to the quality of the discourse in the comments. Here we take a closer look at how discourse character such as quality and anonymity is implicated with reading behavior for users and reporting behavior for journalists.

Readers’ Perceptions of Comment Quality

Most respondents read the comments to some degree (65% read “all the time” or “often”); there was a minority (10%) who reported rarely reading the comments. We examined whether more frequent readers of comments on the site (“all the time” or “often”) evaluate the comments differently than infrequent (“hardly ever” or “never”) readers. Among the frequent readers, the mean agreement with the statement, “I find some of the comments on SacBee.com to be offensive” was 4.83 out of 7 whereas for infrequent readers the mean was substantially higher, 6.31 out of 7. The difference is statistically significant ($t(389) = 4.46, p < .001$).

The reasons people cited for not wanting to read comments included content complaints such as off-topic, predictable, argumentative, or too political as well as that people there were uninformed, mean spirited, uncivil, judgmental, flaming, or making personal attacks. We also looked at the open-ended responses to the question, “Please explain why you do or don’t read comments on SacBee.com” and counted the number of complaints about quality for each level of reading frequency. We found a significant reverse correlation ($r = -.867, p < .02$) between frequency of reading and percentage of responses complaining about quality. For instance, 78% of users who read comments “hardly ever” complained about quality whereas only 5.4% of users who read “all the time” had similar complaints.

Note that despite the correlation between frequency of reading and complaints about quality, we cannot claim that

improving quality would increase reading behavior. For instance, people who read comments “sometimes” may do so for reasons other than quality, but since they only read sometimes and on stories that generate a lot of interest, they may have a skewed perception of the overall quality of the comments on the site. It may be that low quality does lead people to read comments less frequently, but it may also be that more casual, less frequent readers have a biased sampling of the commenting system.

Discourse Character and Journalism

From a journalist’s perspective discourse quality can have a very real and tangible impact on the ability to report a story. Five of the reporters and two of the editors that we spoke to complained about how sources or subjects of stories could be easy targets for criticism or abusive attacks within the comments. In many ways the culture of journalism, reliant on practices of verification and careful sourcing of information, butts heads with the pseudonymous and difficult to tame online space [16]. As the editors put it:

“On the one hand we have this policy of we’re not going to go with anonymous sources and then you have a system in which your sources basically are being attacked by your anonymous commentator” (E7).

“Some people open up and share their lives with us about a particular topic and then they’re attacked” (E1).

In particular, if sources begin to fear the social ramifications of being named in a news article it could make it difficult for reporters to get those sources to speak on the record and consequently to get the information they need for their stories. Reporters have already witnessed these effects,

“I’m starting to have a little bit harder time getting people to agree to be named in a story because they’re concerned about what the commenters are going to say about them once they put their necks out there. It’s a real problem.” (R6)

“If you write a story about welfare recipients and they all go on there and bash the welfare recipient it doesn’t help you get sources down the road.” (R2)

Journalists accept that public figures that are sources or subjects of stories can be and are criticized in the comments, but feel that “normal” people, who are often used to add a voice or to illustrate a particular point of view, should not be subjected to the same level of public criticism. It’s feasible that this attitude stems from the distinction the Supreme Court makes between defamation suits brought against private versus public figures [29]. One reporter questions the implication of this practice for journalism,

“If people want to call people names who are public figures that’s part of the game but are we gonna stop reporting on regular people because

we're gonna have a policy that people can say anything they want?" (R3)

The sentiment in the newsroom by editors and reporters alike was that anonymity was to blame for many of the comment quality problems. Indeed, normative descriptions of journalistic culture espouse the importance of naming sources, most importantly so that consumers of the information can make their own judgment about the credibility of sources [16]. As one editor put it,

"...it doesn't adhere to the principles of the paper itself ... What a beautiful thing it would be if they were not anonymous, so they actually have real people commenting ... We all feel stories are better with real people talking." (E4)

Indeed, journalists felt that the issue of sources being attacked in the comments was exacerbated by the apparent anonymity of the comments. There was the acknowledgment that the paper could get more information from the comments as a result of the anonymity, especially for some kinds of stories (e.g. fires, crime), however, this view was hedged by the argument that there are other anonymous channels already (e.g. telephone), *"The kinds of stuff that people say in comments that is legitimately anonymous can probably be said elsewhere" (R7).*

From the comment writer's perspective, a substantial portion of users (39.3%) believe that anonymity is an important aspect of the system to maintain for their continued usage of it. Of the 298 people who reported that they had commented on the site before, 39.3% indicated that they would *not* be willing to continue posting comments if they had to also provide their real name along with their comment ("cease" group) whereas 60.7% said they would continue posting comments ("continue" group). If we compare the "cease" group with the "continue" group we find that the "continue" group found comments significantly more offensive (agreement with *"I find some of the comments on SacBee.com to be offensive"* $\mu_{\text{continue}} = 5.34$ vs $\mu_{\text{cease}} = 4.74$, $t(297) = 1.94$, $p < .05$). In other words, people who are willing to forego anonymity in commenting are

more likely to find comments offensive, perhaps because they believe less anonymity will improve quality.

In this section we have characterized the implications of low quality online comments in the context of journalism, including impacts on comment reading behavior as well as on the information gathering behavior of journalists. Furthermore, we show that although anonymity is a frequently cited source of concern from the newsroom, it is perceived as an important facet of the system to maintain for a substantial group of users. In the following section we go deeper into the value that readers and writers find in the comments including their motives for use of the system.

Uses and Gratifications of News Commenting

Here we consider the value that users hope to gain from both reading and writing comments. In our survey we asked users how they decide against or in favor of commenting on a story, as well as why they did or did not read comments. Survey responses were analyzed using grounded theory to identify different categories of activity.

Although our analysis allowed for users' reading and writing motives to emerge from the data, we found that the uses and gratifications framework [20, 26] provides a sound structure for our observations. The uses and gratifications framework describes how and why active media consumers engage media in order to satisfy specific needs. It offers a high-level typology of gratifications that people typically seek from media including *information, personal identity, integration and social interaction, and entertainment* [20]. What's unique about the study of commenting systems is that users are in the position not only to obtain gratification through reading comments, but also potentially through *writing* comments. Our goal is to lend insight into how perceptions of comment quality can be addressed by considering individually varying motivations for use. Table 1 shows a summary of user motivations for reading and writing comments including the magnitudes of response.

Comment Writing

Information centric motives involved things such as educating others, answering or asking questions, adding

Writing Motives	Reading Motives
Information (55) Educating others, answering / asking questions, adding information, sharing experience, clarifying, note missing information, balance discussion, correct inaccuracies, factual errors, or misinformation	Information (21) Get more information on story, get additional reporting including updates to story
Personal Identity (59) Expressing an intense emotion or opinion	Personal Identity (10) Validating or comparing personal opinion against community opinion
Entertainment (5) Inject humor into discussion, debate	Entertainment (18) Funny
Social Interaction (24) See reaction of community, persuade others, sympathize, leave condolences, applaud goodness, debate	Social Interaction (102) See perspectives or views from community, see people's true feelings on topic, gauge political response or agenda, take pulse of community

Table 1. Summary of uses and gratifications for reading and writing comments including category counts in parentheses.

information (e.g. facts, insights, background, observations, links, context), sharing first hand experience and personal knowledge, clarifying points, noting missing information, balancing a discussion, or pointing out incorrect information (e.g. inaccuracies, false statements, factual errors, or misinformation). Many of these information centric motives can be considered “social” in the sense that they add concretely to the public discourse.

Personal identity motives manifested as people wanting to express some sentiment (e.g. a feeling such as anger or outrage, an opinion, or to vent). People often reported that it was often an intense interest or a passionate emotional response to a story or issue that would initiate the desire to comment. *Entertainment* motives were fairly straightforward: some people enjoyed injecting humor into the discussion or found it enjoyable to debate with others.

Social interaction based motives for writing included affective uses such as the desire to sympathize, leave condolences, or applaud good work by reporters or other commentators. These affective uses are social in that they are emotional expressions explicitly directed at other people in the online community. People would also sometimes comment in order to see the reaction of other commentators, in effect comparing others’ responses to their own. Another motive was to persuade others, such as trying to get the newspaper to take some action or cover a particular story.

Reasons for *not* wanting to comment included issues of privacy and vulnerability. Pseudonyms can often be traced back to real people through IP addresses (or in some cases through unique names included in the handle). Two percent of respondents stated that such incomplete anonymity was a reason for not wanting to write a comment. Some people also did not contribute because they thought their contribution would not matter or would be overlooked.

Comment Reading

Information-based motives included people’s desire to get more information such as getting updates to a story or getting additional reporting or information via the comments. In contrast to the informational writing motives which served to disseminate and share information, the informational reading motives focused on the individual *obtaining* information from the comments.

Personal identity motives involved validating or comparing one’s opinion against that of the community. For instance, someone might be interested in seeing if their own opinion was already expressed in the comments or in how the opinions expressed fit with their own views. Some people also mentioned that *entertainment* was a factor in their choice to read the comments.

Reading motives relating to *social interaction* included people wanting to see a divergence of other views in the comments such as other perspectives or dimensions, to see people’s true feelings, or to get the community’s current thinking on an issue. People were also interested in gauging

the political response to a story, taking the pulse of the community (including the public or local reaction and popular perception) or to see political agendas.

Motives and Perceptions of Quality

A useful distinction made by the uses and gratification framework is that between *gratifications sought* and *gratifications obtained* [20, 26]. This distinction highlights the potential mismatch between what an individual wants from a media experience and what they actually end up getting from it. Furthermore, personal temperaments and differences may influence the wants and needs that an individual has at any given time. To assess how differences in motivation might affect perceptions of comment quality, we compared responses of users solely motivated for *individual centric* reasons (information, personal identity, and entertainment motives) and users that solely mentioned motivations in the *social interaction* category. The mean agreement with “*I find some of the comments on SacBee.com to be offensive*” was 5.62 out of 7 for *individual-centric* people, whereas the mean for people motivated for social interaction reasons was lower, 4.8 out of 7 ($t(122) = 1.94, p = 0.05$). This finding suggests that differences in reading motivations can influence perceptions of quality and offensiveness in the comments, with people motivated by social interaction reasons less offended by the comments on average. These differences may indicate that some of the issues in perceptions of quality stem from a mismatch between motives and expectations for the system.

Strategies for Addressing Comment Quality

We have demonstrated that comment quality is a significant issue for the people participating in the online space and affects how journalists perform their work. Here we consider a number of approaches undertaken at SacBee.com towards improving quality. These approaches include *flagging*, *moderation*, and *engagement*. In particular, we will consider the effectiveness of these approaches, issues in their deployment, and how their implementation aligns with the values of the newsroom.

Flagging

SacBee.com asks users to flag comments as abusive if they do not adhere to the commenting policy. With enough flags, a comment will be automatically removed from the site, or at minimum will get sent to the moderation queue for review by a moderator. To uncover the reasons and context of why and under what circumstances users flag comments, we asked users in the survey about their flagging behavior.

Many of the users who said they use comment flagging on the site do so for all of the reasons outlined in the commenting policy including: threats, hate speech, offensive language, ad hominem attacks, off-topic material, and spam. However, there are some notable other reasons mentioned by users for flagging comments as “abusive”, including: legally inspired reasons such as libel or privacy, disparaging or distasteful remarks about the dead, and more personal reasons like revenge and disagreement or to deal with ignorant posters or inaccurate facts.

There were also a number of reasons readers gave for why they had not or did not want to flag. 6.4% of respondents felt that it was not their position to be adjudicating when someone could or couldn't say what they wanted. They felt that everyone is entitled to an opinion and sometimes referred to the First Amendment of the U.S. constitution as the basis for this rationale. There were three respondents who did not want to flag because they recognized that it was a subjective process with blurry lines that might not correspond to their own notions of appropriateness.

Editors are well aware of the issues with the flagging approach. One editor attributed some of the problems with the flagging system to the diversity of the community, "*We haven't been able to create a community that self regulates because it's not a community in the sense that everybody has the same interests at heart.*" And although we cannot make any claims about the magnitude of the abuse, moderators and editors from the newsroom did corroborate some of the more deviant uses of the flagging system:

"I guess the ones where you have to be fairly discerning is when you see that certain viewpoints whether by one person or a group of people continue to be flagged and you read into them and you find out well, these are being flagged by somebody why disagrees with them and just wants their point of view out" (M6).

Moderation

As described before, 17.2% of comments on the site were either marked as hidden or were deleted from the site via the moderation system. Moderators acknowledged that *most* comments on the site do not pose a problem, but that the topic of a story often impacts the tone of the response in the comments. Indeed, studies have shown a relationship between topicality and perceived politeness in online discussion groups [2]. In the newspaper context, controversial topics such as immigration, gay marriage, crime, and welfare can easily generate such a large response as to overwhelm moderators – even when only considering comments that had already been flagged by the community.

Outsourcing the moderation process is one avenue for dealing with an increase in comment volume. But 60% of editors and reporters expressed reasons for not wanting to outsource abuse moderation. Editors felt that editorial control was necessary for doing a good job moderating since comments are seen as a form of content. Editors also thought it would be important to know the context, the issue, the background, and the geographic area to do an adequate job, all of which would be difficult for an external and potentially distant company to learn. As one editor said,

"I think that moderating the comments is an editorial decision and needs to be made by an editorial obligation...and to know not only the newsroom but the community." (E2).

Reporters were likewise concerned about outsiders not having the same journalistic standards, or not being able to interpret local issues in a way consistent with what would be expected by a Sacramentan:

"It seems like you'd run the risk of an outsider interpreting what would be inappropriate for people who live somewhere else...someone in Iowa or India or wherever may not understand the degree of emotion that we might tolerate here over that issue versus what they would tolerate in their community"(R4).

"Any two people are going to have different judgment about what to put on the web and what not to. I think the closer they are to the environment that they're dealing with the better." (R5)

Indeed, even among their own newsroom, editors and moderators are well aware of the subjective nature of their task. They all acknowledge that there are sometimes gray areas and that different moderators will have different standards that they apply to the same material:

"We look for profanity, obscenity, things that are vulgar, things that are personal attacks, and it's all subject to each individual's assessment of what's a personal attack." (E1)

"What one person thinks racially prejudiced, another may think is a criticism of culture. That's usually the toughest question" (M3)

Despite the acceptance of the subjective nature of the task of moderation, our findings indicate the newsroom culture was reluctant to release control of content moderation to "outsiders" who might not have the same editorial standards or locally meaningful understandings of issues to make the same kinds of hard subjective decisions that they do. While consistent with traditional norms of content control and gatekeeping in mass media, these findings highlight the tension between a desire to have better quality comments and a reluctance to accept outside help in editorial tasks.

Engagement by the Newsroom

Of the survey respondents who read comments "All the time" or "often" the mean agreement with the statement "*I would like reporters to respond to the comments on their stories in order to clarify or answer questions*" was 5.79 out of 7. In an open ended question about other things that might improve the experience of commenting, the idea of engagement was elaborated to include having reporters or editors respond to false information, clarify information, highlight good comments, or create competition to have one's comments highlighted.

Almost all of the reporters and editors (13 of 15) do engage their readership to varying degrees via email, but are less inclined to do so via comments. The idea that this engagement can soothe the commenter is understood by

many reporters: *“Sometimes I get very heated angry emails and if I respond to them sometimes people say ‘oh I didn’t know that.’ And it changes their tone”* (R3). The major difference then with the comments is that they are public and persistent records of the engagement. There was some confusion in the newsroom about whether or not reporters were even permitted to comment on stories – many believed that they were not allowed to. This situation in turn has led to some tension for reporters who believe that more engagement with the readership is necessary: *“Part of the problem is that we think we can just let the comments spin out there without our touching them”* (R6).

Many of the reporters (5 of 7) indicated that they would want to engage readers more via the comments, or at least that it would be a valuable experiment to try out:

“If it took 40 minutes to an hour out of your day I think it would be an interesting experiment to see what happens when they know that the writers are still around and you would get a different level of interactivity and possibly connection with the author of the story.” (R2)

“There have been many cases where a reader will raise a question about the story that the story itself didn’t answer ... and there’s no way for a reporter to go into the comments and submit a comment that responds to that ... I think there ought to be some formal and recognized system for the staff of the newspaper to respond to reader comments. To make it more of a dialogue, to correct inaccuracies that readers sometimes create when they submit comments.” (R4)

“I think that we should be able to comment to correct factual information...or answer questions that can be objectively answered.” (R7)

Despite the predominant reader and reporter perspectives on engagement, there was still pushback from editors in terms of allowing more newsroom engagement in the comments. Some editors for instance, thought that the comments were the purview of the users of the site and that the newsroom staff should not respond. Another pragmatic concern was that it would take too much time away from other tasks. Finally, some editors thought that reporters might get overexposed in the comments. One editor put it thus,

“It gets into a dangerous area with reporters of we’re trying to maintain our standards of accuracy, fairness, and credibility. And once they start weighing in to a dialogue, it’s very easy to move over that fence into talking directly to a commenter and then it’d be difficult not to have opinions. And I don’t want them to go there.” (E4)

At the same time, several editors acknowledged that for more objective questions or issues of accuracy, having reporters respond in the comments could be very beneficial:

“I think it helps when the reporter answers questions or weighs in verifying facts or even getting at the intent question.” (E6)

“I wouldn’t have a problem at all if there was a commenting thread that was inaccurate if a reporter just posted something pretty basic, just very fact based ... just to clarify the thread.” (E5)

In terms of reporter engagement, the editors’ fears relate to how objectively reporters can engage the commenters. If reporters expose their personal opinions in the comments this could lead to accusations of bias or to a loss of credibility for the paper. The unmediated nature of the comments also made at least one editor uncomfortable since he didn’t have the ability to edit or oversee his reporters.

Many reporters, columnists, and editors do currently engage the comments via *reading* as long as it’s within reasonable time limits. Often this means they might not get past the first 20-30 comments on a story (granted, 88% of stories in our sample got less than 30 comments). Reporters or columnists writing about more controversial topics may need more time or the help of an editor in order to give the comments a fair read. Journalists found value from reading the comments in many ways including finding tips and sources for current or future stories, getting a feel for the sentiment response to a story, or for receiving criticism. Many journalists also expressed that they engage in flagging of abusive comments, though inconsistently since they spend variable amounts of time reading the comments.

DISCUSSION AND IMPLICATIONS FOR DESIGN

Consistent with prior studies of professional inertia in newsroom innovation [7, 25], we found there to be a number of cultural and organizational barriers to implementing strategies for quality improvement. Here we consider implications for designing interventions that acknowledge this complex organizational milieu.

Our results indicate that, on an individual level, there can be mismatches between the benefits sought from the commenting system and the benefits achieved. While the motives for both writing comments as well as for reading them seem to dovetail fairly well, our data suggests that individual differences in reading motivations may have an impact on quality perceptions. Matching these motives to customized views of the commenting system may help address some perceptions of poor quality.

A technical approach that could help alleviate the individual expectations for uses and gratifications of reading comments could involve building advanced filtering tools that could be used to sort and organize comments based on the sentiment (i.e. subjectivity or objectivity) of the words used in a comment [24]. Subjectively categorized comments can furthermore be separated into positive and negative responses [10]. For instance, objective messages might appeal more to those who are motivated to obtain information, whereas subjective messages may appeal more

to those interested in social interaction, community pulse, or entertainment. Another avenue for ranking comments could be based on the politeness of comments so as to allow readers more latitude in the tone of the comments they see [2]. The same sorting tools built for readers could also help journalists extract more value from the comments.

The flagging procedure used on the site is largely effective for identifying abusive comments. However, there are also vindictive uses of flagging which increase the workload of moderators who have to sift through the bogus flags. Abuses of the flagging system could be addressed by developing both a policy of warning people who are suspected of bogus flagging as well as moderation tools to track users' flagging behavior. Such a policy would be a nominal extension to the existing moderation system. A user marked as a flag abuser could be dealt with by simply not recording any flags from that person in the future (i.e., a "bozo filter" for flags). This method would have the effect of lowering the overall volume of flagged comments, while increasing consistency and trustworthiness of the remaining flags.

Qualitative review of the survey responses concerning flagging did uncover some legitimate uses for flagging such as for marking disparaging remarks about deceased people, or for concerns over libel or privacy. Currently there is no way for people to explicitly tag comments as such in the process of flagging. Providing extra tags for flags may raise awareness for relevant issues and channel those concerns to someone with legal training within the newsroom.

There is a concern in the newsroom that certain story topics arouse so much interest in commenting that they can easily overwhelm the moderators. Calls for commenting bans on problematic local topics yield free speech and censorship concerns for users and journalists alike. Using a database of moderated (i.e. flagged or deleted) comments it may be possible to learn text models to help predict high-volume or problematic stories based on story topicality. This in turn would enable editors to, at the very least, set expectations or work shifts for moderators on any given news day.

One organizational approach toward dealing with the growing volume of comment moderation involves outsourcing. However, this idea was rejected by almost all the journalists we spoke to on the grounds of it being a subjective editorial process that was hard to keep consistent even *within* the organization. Considering the importance of editorial control in making a *final* moderation decision a compromise would entail augmenting community flagging activity with outsourced professional flagging activity. Moderators in the newsroom would still have the final say in what comments were removed or retained but the outsourced flagging would be a consistent initial filter.

A more tenable approach toward quality improvement is to have reporters engage readers in the comments more directly. By having reporters read and respond to comments this would also increase the number of trusted flags in the system. Clear standards would need to be developed about

when it is or is not appropriate for a reporter to engage the comments. For instance, factual corrections, or factual questions may be alright whereas something about editorial decisions or other subjective components may be off limits. The barriers to having this approach accepted by editors could be addressed by introducing better facilities for editors to *monitor* how their reporters are engaging readers in the comments. If, for instance, a reporter got into a heated exchange with a reader, the editor could step in to mediate or calm the situation down. Ideally, such a monitoring system would form the basis for an editor's dashboard so that editors could be confident that the editorial content that reporters were generating in the comments was consistent with professional and institutional standards.

CONCLUSIONS AND FUTURE WORK

Online news commenting systems present a highly visible, unique, and important challenge for information in local and online communities. The unique aspects of these communities include often-volatile participation patterns, imbalance between professional and amateur content, and interaction between regular users and other actors in various official capacities. News organizations are tasked with providing forums which engender constructive discussions which coalesce and strengthen their online communities. In this work we have (1) shown the extent of the impact of low quality comments on both users and journalists, (2) explicated how individual differences in reading motivations can impact perceptions of quality, and (3) considered aspects of flagging, moderation, and engagement as strategies for improving quality. From these results we derive value-sensitive design implications for improving the discourse quality in online news commenting systems.

Comment organization, sorting, and filtering tools based on content analysis (e.g. sentiment) are promising avenues for structuring the commenting experience and helping meet users' expectations and get the most individualized value out of comments. We would like to implement these features, as well as relevancy and novelty detectors [5, 9] and evaluate the impact they have on the user experience of comment reading. While there has been some work in applying automatic content analysis to classify or score message quality [8, 32], there is a dearth of evaluation examining how these techniques are perceived and used by people. For a naturalistic evaluation we would be interested to see if an improvement in the perceived quality of comments seen on a news story also increases the sense of community or frequency of usage. Furthermore, many of these automatic tools could also be beneficial to help journalists extract value from the comments, such as to summarize the discussion and perspectives there [5, 22].

There are a number of future analyses we intend to carry out on our dataset of 54,540 comments. Our results indicate an even greater proportion of singleton posters (46.5%) than studies of, for instance, USENET [33]. It would be interesting to examine the relationship between unequal distribution of comment participation on the tone and

quality of discussion across different stories. What is the relationship between comment writing frequency and the subjectivity or objectivity of comment content? How much does topicality really affect the quality of comments across different news stories? Are there temporal patterns to more or less quality comments? Can we extract and utilize the social connections or community affiliations (if any) among users who comment on different topics? Furthermore, is there a connection between comment quality and the “localness” [12] of the person contributing? A deeper understanding of the relationships between topicality, comment content, and social relationships may lead to further insights in how to manage and improve the quality of comments specific to news sites.

ACKNOWLEDGMENTS

We wish to thank the Sacramento Bee for their support and cooperation. The first author also acknowledges the support of the AAAS via the Mass Media Fellowship program, the University of Bergen Department of Information Science and Media Studies, and the CRA for a Computing Innovation Fellowship (CIF-197).

REFERENCES

- Bergström, A. The Reluctant Audience: Online Participation in the Swedish Journalistic Context. *Westminster Papers in Communication and Culture*, 5 (2). 60-80.
- Burke, M. and Kraut, R., Mind Your Ps and Qs: The Impact of Politeness and Rudeness in Online Communities. in *Computer Supported Cooperative Work (CSCW)*, (2008).
- Chung, D.S. Profits and Perils: Online News Producers' Perceptions of Interactivity and Uses of Interactive Features. *Convergence*, 13 (43).
- Cosley, D., Frankowski, D., Kiesler, S., Terveen, L. and Riedl, J., How Oversight Improves Member-Maintained Communities. in *Proceedings of CHI*, (2005), 11-20.
- Diakopoulos, N., Naaman, M. and Kivran-Swaine, F., Diamonds in the Rough: Social Media Visual Analytics for Journalistic Inquiry. in *IEEE Conference on Visual Analytics Science and Technology (VAST)*, (2010).
- Díaz Noci, J., Domingo, D., Masip, P., Lluís Micó, J. and Ruiz, C., Comments in News, Democracy Booster or Journalistic Nightmare: Assessing the Quality and Dynamics of Citizen Debates in Catalan Online Newspapers. in *International Symposium on Online Journalism*, (2010).
- Domingo, D. Interactivity in The Daily Routines of Online Newsrooms: Dealing with An Uncomfortable Myth. *Journal of Computer-Mediated Communication*, 13. 680-704.
- Fortuna, B., Rodrigues, E.M. and Milic-Frayling, N., Improving the Classification of Newsgroup Messages through Social Network Analysis. in *Conference on Information and Knowledge Management*, (2007).
- Gilbert, E. and Karahalios, K., Understanding Deja Reviewers. in *Computer Supported Cooperative Work (CSCW)*, (2010).
- Gindl, S. and Liegl, J., Evaluation of Different Sentiment Detection Methods for Polarity Classification on Web-Based Reviews. in *Computational Aspects of Affectual and Emotional Interaction*, (2008).
- Gurzick, D., White, K., Lutters, W. and Boot, L., A View from Mount Olympus: The Impact of Activity Tracking Tools on the Character and Practice of Moderation. in *GROUP*, (2009).
- Hecht, B. and Gergle, D., On the "Localness" of User-Generated Content. in *Computer Supported Cooperative Work (CSCW)*, (2010).
- Hermida, A. and Thurman, N. A Clash of cultures: The integration of user-generated content within professional journalistic frameworks at British newspaper websites. *Journalism Practice*, 2 (3). 343-356.
- Kiesler, S., Siegel, J. and McGuire, T. Social Psychological Aspects of Computer-Mediated Communication. *American Psychologist*, 39 (10). 1123-1134.
- Kilner, P. and Hoadley, C. Anonymity Options and Professional Participation in an Online Community of Practice *Computer Support for Collaborative Learning*, 2005.
- Kovach, B. and Rosenstiel, T. *The Elements of Journalism: What Newspeople Should Know and the Public Should Expect*. Three Rivers Press, 2007.
- Lampe, C., Johnston, E. and Resnick, P., Follow the Reader: Filtering Comments on Slashdot. in *Proceedings of CHI*, (2007), 1253-1262.
- Lampe, C. and Resnick, P., Slash(dot) and burn: distributed moderation in a large online conversation space. in *Proceedings of CHI*, (2004), 543-550.
- Manosevitch, E. and Walker, D. Reader Comments to Online Opinion Journalism: A Space of Public Deliberation. *Online Journalism Symposium*, 2009.
- Miller, K. *Communication Theories: Perspectives, Processes, and Contexts*. McGraw-Hill, 2004.
- Nagar, N.a. The Loud Public: Users' Comments and the Online News Media. *Online Journalism Symposium*, 2009.
- Nam, K. and Ackerman, M.S., Arkose: Reusing Informal Information from Online Discussions. in *GROUP*, (2007).
- Otterbacher, J., "Helpfulness" in Online Communities: A Measure of Message Quality. in *Proceedings of CHI*, (2009).
- Pang, B. and Lee, L. *Opinion Mining and Sentiment Analysis*, 2008.
- Paulussen, S. and Ugille, P. User Generated Content in the Newsroom: Professional and Organisational Constraints on Participatory Journalism. *Westminster Papers in Communication and Culture*, 5 (2). 24-41.
- Perry, D.K. *Theory and Research in Mass Communication: Contexts and Consequences (2nd Edition)*. Lawrence Erlbaum Associates, 2002.
- Powazek, D. *Design for Community: The Art of Connecting Real People in Virtual Places*. Waite Group Press, 2001.
- Purcell, K., Rainie, L., Mitchell, A., Rosenstiel, T. and Olmstead, K. Understanding the participatory news consumer *Pew Internet*, 2010.
- Solove, D.J. *The Future of Reputation: gossip, rumor, and privacy on the internet*. Yale University Press, 2007.
- Strong, D., Lee, Y. and Wang, R. Data Quality in Context. *Communications of the ACM*, 40 (5). 103-110.
- Taylor, R. *Value-Added Processes in Information Systems*. Ablex, 1986.
- Wanas, N., El-Saban, M., Ashour, H. and Ammar, W., Automatic Scoring of Online Discussion Posts. in *Workshop on Information Credibility on the Web (WICOW)*, (2008).
- Whittaker, S., Terveen, L., Hill, W. and Cherny, L., The Dynamics of Mass Interaction. in *Conference on Computer Supported Cooperative Work (CSCW)*, (1998).
- Wise, K., Hamman, B. and Thorson, K. Moderation, Response Rate, and Message Interactivity: Features of Online Communities and Their Effects on Intent to Participate. *Journal of Computer-Mediated Communication*, 12 (1).