Systems for Improving Online Discussion

Amy X. Zhang
MIT CSAIL
Cambridge, MA 02139 US
axz@mit.edu

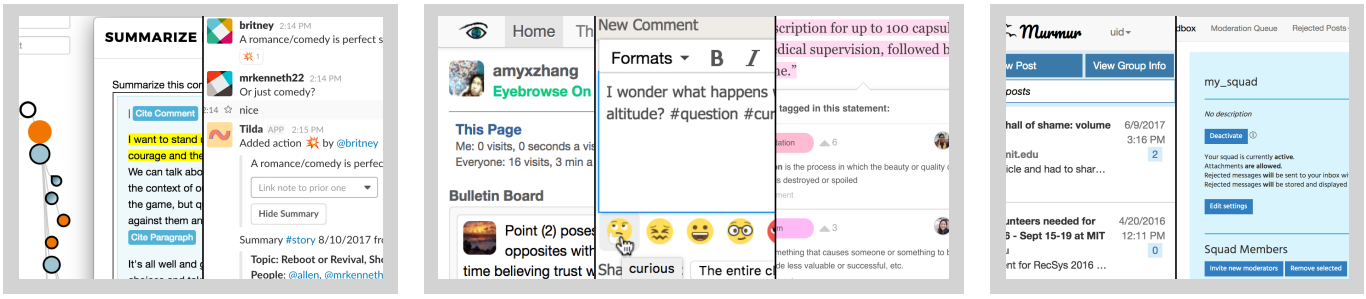

Figure 1: Novel online discussion systems. The first set of systems helps manage large discussions, while the second provides annotation tools to understand context, and the third allows users to have fine-grained control over the delivery of messages.

ABSTRACT

More and more of the discussion that happens now takes place on the web, whether it be for work, communities of interest, political and civic discourse, or education. However, little has changed in the design of online discussion systems, such as email, forums, and chat, in the decades they have been available, even as discussions grow in size and scope. As a result, online communities continue to struggle with issues stemming from growing participation, a diversity of participants, and new application domains. To solve these problems, my research is on *building novel online discussion systems* that give members of a community direct control over their experiences and information within these systems. Specifically, I focus on: 1) tools to make sense of large discussions and extract usable knowledge from them, 2) tools to situate conversations in the context what is being discussed, as well as 3) tools to give users more fine-grained control over the delivery of content, so that messages only go to those who wish to receive it.

Author Keywords

online discussion; social media; threaded discussion; synchronous communication; summarization; annotation; moderation; deliberation; online communities

UIST '17 Adjunct, October 22–25, 2017, Quebec City, QC, Canada.

ACM ISBN 978-1-4503-5419-6/17/10 $15.00.
https://doi.org/10.1145/3131785.3131845

ACM Classification Keywords

H.5.3. Group and Organization Interfaces: Web-based interaction

INTRODUCTION

Discussions on systems such as email, forums, and chat have been pervasive on the internet since its inception. They contain a diversity of rich information and experiences, including differing opinions on an issue, anecdotes, humor, explanations, coordination, and deliberation. Over the course of thousands of comments, even open mathematics problems on Polymath can be solved and controversial edits on Wikipedia settled. However, online discussion tools are still remarkably primitive, barely changed from their origins in forums, email, and instant messaging. As a result, problems with discussion systems persist and are now exacerbated by growing participation in online discussion and new discussion domains.

Too often, discussions online lead to polarizing interactions or peter out with no resolution, wasting valuable time and effort between participants. Sometimes this happens because of the *scale of the discussion*, making it difficult for participants to get an overview of what happened or refer back to specific content. A related problem is the *overload of disparate discussions* that are all lumped into the same place with little context of what the conversation is about. Other times the problems are due to the *diversity of participants*, leading to tensions stemming from different expectations of appropriate behavior, modes of access, and opposing perspectives. In the worst cases, this can escalate into harassment of individuals as well as an overall decline in civility [1].

In my thesis, I introduce a diverse set of systems that reimagine outdated discussion designs to mitigate the problems that

discussion systems face. These systems confront the growing pains of discussion systems by providing tools to users that give them greater control over their experiences and information online.

- The first set of systems help make sense of large discussions by providing collective *summarization* capabilities within discussion systems, such as the ability to bridge between forums and wikis using recursive summarization (Wikum), as well as to organize synchronous chat streams into skimmable, linked notes (Tilda). These systems for summarization of discussion both provide a mechanism for people to deal with scale as well as a way for people to reflect and collaborate on a holistic understanding of what was said.

- The second set of systems situate conversations in the proper context by providing *annotation* capabilities when having discussions, such as the ability to embed discussions as annotations within primary content such as webpages (Eyebrowse) or textbooks (Nota Bene), or mark up articles and discussions with tags characterizing their framing (Pano). By placing the conversations where their annotation resides, these systems allow conversations to be viewed in context instead of in a single overloaded space. Also, the use of annotations to label source materials can help give high-level, visual signals to readers about the content and where to dive in.

- The third set of systems provide tools that give users *finer-grained control over the delivery of content*, including what messages they see, who in turn sees their messages, as well as how the messages get accessed and delivered. This is achieved through giving users a richer vocabulary of actions they can take to collectively and individually filter and direct conversations in mailing lists (Murmur), as well as giving users the ability to recruit friend moderators to protect them from email harassment (Squadbox). These two systems help tackle issues surrounding how to accommodate the diversity of norms in a community.

THESIS STATEMENT

We can build tools to enable users and communities to summarize, annotate, and more precisely control the delivery of their discussions, and the use of these tools will improve the experience and outcomes of online discussions.

SYSTEMS

Making Sense of Large Discussions

The first set of systems tackle how large discussions can be difficult for getting an overview. Here, a short textual *summary* is the traditional solution, but this is time-consuming for a single person to complete. To address this problem, we consider how a group of people could individually contribute small amounts of work to refine a large discussion into a dynamic textual summary, called a *summary tree* that can be explored at varying levels of detail. We design a workflow for creating a summary tree using the idea of *recursive summarization* of a discussion, where users build summaries of small sections of the discussion, small sets of those summaries are then aggregated and summarized, and so on, until the entire discussion is summarized. To explore the design space of this process, we present **Wikum**[1] [13], a system for creating wiki summaries recursively and reading a discussion overlaid with summaries. From lab evaluations, we see that users found Wikum easier to use and faster for writing summaries than a Google Docs file, mimicking a wiki editing environment. We also find that readers were better able to skim over discussions using a summary tree rather than a word document.

We are conducting additional studies and development to further explore ideas around collaborative summarization. First, we have established a collaboration with the **Wikimedia Foundation** to study Wikum's potential use as a tool for Wikipedia editors to summarize deliberative discussions on their Talk Pages. From interviews with Wikipedia dispute resolvers, we are adding features to more easily determine when a single person is dominating the conversation or using multiple "sock-puppet" accounts to create the appearance of consensus. In addition, we are interested in exploring how participating in summarization while in the process of conversing or reading can produce more reasoned debate and nuanced perspectives. We are planning additional features for Wikum to support these studies, including summarization of ongoing discussion in addition to static discussion, allowing "meta" discussion around the summary tree artifact, and allowing summarizers to create different user roles and permissions within the tool.

We are also conducting development of **Tilda**, a tool developed at Microsoft Research that converts long chat conversations into short summaries that link back to the original discussion. The tool works by coordinating inputs from discussion participants to organize an ongoing stream of synchronous communication into structured summaries of discrete conversations. This allows users to easily get an overview of the chat for when they've missed a batch of messages, are newcomers to the community, or need to search through old chat logs. From interviews, we discover what information users want to glean from chat as well as what format is preferable for the summaries. As part of this work, we plan to collect and release a public dataset of annotated and summarized chat logs towards building machine learning models to support the summarization process.

In the future, I am interested in exploring how humans and automatic techniques can work in tandem to construct summaries. Part of this work involves examining the summarization process to see how it can be broken down into discrete steps that can be helped by automatic techniques. Another aspect includes understanding empirically what humans desire within summaries and centering evaluations around the needs of humans in particular contexts. For instance, in user studies of Tilda, we find that users often prefer summaries structured by high-level discourse acts as opposed to unstructured abstractive or extractive sentences. The characterization of different types of discussion using common discourse sequences [11] may also be helpful towards understanding how

[1] `wikum.csail.mit.edu`

summarization needs and workflows for creating summaries change depending on the type of conversation.

Situating Conversations in Context

Even in cases where individual discussions themselves are not very long, when there are many threads to go through, understanding what to read or respond to can be difficult. One reason for this issue is that, traditionally, all conversations are placed in a single forum or inbox regardless of their content.

To combat this issue, we develop annotation systems that support discussions that live in the "margins" of webpages or textbooks. One such system is **Eyebrowse**[2] [10], a browser extension and web tool that allows users to share aspects of their browsing history, such as when they are active on particular sites. This allows people to "bump into each other" in spaces on the web and have serendipitous live conversations or asynchronous discussions that are then tied to the pages they are on. From a week-long field study of Eyebrowse within 4 friend groups, we found that people were interested in sharing aspects of their browsing more publicly, and some of the groups found the conversational aspects of Eyebrowse engaging. There were several instances of participants having serendipitous interactions around the content of a page they were both on, such as discovering that two participants were attending the same conference. Moving forward, we aim to further explore the use cases around within page chat on the web, such as for news readers, online learners, or developers.

A second system is **Nota Bene (NB)**[3] [14], a tool for students to have discussions in the "margins" of online textbooks anchored to a specific passage on the page. This allows students to read and concurrently view discussions in context of what is being discussed. In my work related to NB, I consider how user signals such as emojis and hashtags in conversation could provide signals for readers about what discussions to focus on and what reactions a particular passage elicits [12]. This could also help instructors cut through the noise of discussion, especially in large classrooms such as MOOCs, to find unanswered questions or resolve disputes, or distinguish interesting from confusing passages for future revision.

Finally, research into political discourse demonstrates that participants often "talk past each other" when arguing their point of view due to a lack of context about the other side's moral values. We build a system called **Pano**, built on top of our Eyebrowse framework, that allows people to tag comments and articles with moral framing in order to become aware of their own and their peers' preferred frames. The framing annotations are also used to differentiate between conversations and news articles of different frames within recommendations in the Eyebrowse extension. A 10-day field study we conducted provides preliminary evidence that users' understanding of differing moral values improved after the use of our tool and also allows them to better frame their own arguments in the values of the opposing side.

In addition to work on Pano, I would like to continue to develop systems to promote reflection and thoughtful, evidence-based discussion along with tools to help users engage in meaningful dialogue. Related to this line of work is understanding how to curb detrimental discussion such as the spread of misinformation or exaggerated claims in social media. Building on ideas from Eyebrowse, a tool could allow a user to give trusted individuals or the crowd the power to annotate content on their social feed with refutations, sources, or alternative content, in order to limit the spread of misinformation, while in the process teaching media literacy to annotators. We have developed an initial prototype of one instantiation of such a tool, called Baitless, that allows the crowd to rewrite the headlines of articles that appear in an RSS feed in order to combat "clickbait" headlines.

Finer-Grained Control of Delivery

A separate problem that discussion communities face is contested norms due to the diversity of participants. From studying existing group communication systems such as mailing lists, Facebook Groups, and forums to find their perceived pros and cons [9], we notice varied perspectives as to how members of the same mailing list should access and post to the list, leading to tensions, as well as a general hesitancy towards posting for fear of spamming.

One solution for this problem is to give users more fine-grained control over all aspects of the delivery of content, so that all the members of a community are working together to ensure that content gets delivered only to those who want to receive it. Drawing from our exploration into tensions within group communication, we present **Murmur**[4], a mailing list system that aims to keep the benefits of email, such as greater confidence that messages will be seen, while introducing new features that are present in more modern systems such as Facebook, such as social moderation. Rather than using algorithmic curation, which puts the delivery of content in the hands of a model, Murmur allows users to have more explicit and fine-grained control to filter, block, follow, and otherwise curate how and to whom discussions are sent and received.

We plan to extend Murmur's flexibility by developing a novel constrained language and corresponding graphical user interface to allow senders and receivers to construct their own rules about how messages should propagate. We also plan to make Murmur a bridge between different discussion platforms such as email, forums, and Facebook Groups, so that members of a group can choose what interface to use.

As part of our development of Murmur, we also study how novel moderation capabilities within communities can curb problems such as harassment and the violation of norms in online communities. These additional moderation capabilities have led to the creation of **Squadbox**[5], a tool for people facing harassment in their personal inboxes to recruit their friends as moderators. In this project, we seek to design ways for moderators and targets of harassment to work together to effectively combat harassment. From interviews with over 10

[2] `eyebrowse.csail.mit.edu`
[3] `nb.mit.edu`
[4] `murmur.csail.mit.edu`
[5] `squadbox.csail.mit.edu`

targets of social media and email harassment, we have learned how users wish to interface with moderators to personalize their strategy for dealing with harassment. We are now planning user studies to understand to what extent hand-crafted filters and human-in-the-loop models can help with moderation and what circumstances need delegation to moderators. Finally, we will be investigating the benefits and tradeoffs of having friend, paid stranger, or volunteer moderators.

RELATED WORK

To reduce scale and manage what content is delivered, many discussion interfaces today have some form of filtering, such as collective social moderation using voting. Many sites and researchers have also experimented with automated filtering, such as for detecting spam and trolls [1]. However, these filters can be opaque and set community-wide, while Murmur gives a user more personalized and rich set of levers beyond simple inputs such as upvotes. Additionally, filtering can only go so far as they may block minority opinions and also may lead to "filter bubbles" when only one point-of-view is represented [8]. Consensus documents such as Wikum's summary trees can provide a more holistic overview. Other approaches aimed at reducing uncivil or harassing content employ moderators to filter comments or use community mechanisms such as flagging [2]. Systems such as HeartMob allow people to provide support after the fact to others who have experienced harassment. These systems demonstrate how non-automated approaches like friend-moderation within Squadbox can often be necessary to sustain a community.

Other work has focused on novel visualization or presentation of discussions to give an overview. There are researchers that have explored how to convey mood or temporal activity [3]. Other examples include Opinion Space [4] or pol.is, interfaces that map users' comments on a two-dimensional space, ConsiderIt [5], which allows users to build up pro-con lists on different issues, and Arkose [7], which group comments into high level clusters to summarize. However, many of these interfaces do not allow for any actual discussion of the opinions presented, which is necessary for nuanced deliberation. Other work aims to better contextualize conversation such as within annotation systems such as Hypothes.is or Genius. We draw upon these systems in our design of features within Eyebrowse and NB. Finally, the system Reflect [6] focuses on paraphrasing as a means of reflection in online discussions. Our systems, such as Wikum and Pano, build on these ideas for encouraging reflection while also providing useful information to reduce scale or aid recommendation systems.

CONCLUSION

My work focuses on building discussion systems that give users greater power to manage and organize their discussions. Through the use of these tools, communities will be able to better manage overload, resolve tensions, and extract meaningful knowledge in the face of growing participation and more diverse membership. These tools will be crucial going forward as online discussion systems play a larger role in the world due to increased participation in online education, remote work, online political discourse, social media, and online civic deliberations.

ACKNOWLEDGEMENTS

My work is supported by a Google Research Fellowship and the National Science Foundation. I am grateful for all of my collaborators and especially my thesis advisor, David Karger, professor of EECS at MIT CSAIL.

REFERENCES

1. Cheng, J., Danescu-Niculescu-Mizil, C., and Leskovec, J. Antisocial behavior in online discussion communities. In *Ninth International AAAI Conference on Weblogs and Social Media* (2015).
2. Crawford, K., and Gillespie, T. What is a flag for? social media reporting tools and the vocabulary of complaint. *New Media & Society* (2014).
3. Donath, J., Karahalios, K., and Viégas, F. Visualizing conversation. *Journal of Computer-Mediated Communication 4*, 4 (1999), 0–0.
4. Faridani, S., Bitton, E., Ryokai, K., and Goldberg, K. Opinion space: a scalable tool for browsing online comments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2010), 1175–1184.
5. Kriplean, T., Morgan, J., Freelon, D., Borning, A., and Bennett, L. Supporting reflective public thought with considerit. In *CSCW*, ACM (2012), 265–274.
6. Kriplean, T., Toomim, M., Morgan, J., Borning, A., and Ko, A. Is this what you meant?: promoting listening on the web with reflect. In *CHI*, ACM (2012), 1559–1568.
7. Nam, K. K., and Ackerman, M. S. Arkose: reusing informal information from online discussions. In *Proceedings of the 2007 International ACM Conference on Supporting Group Work*, ACM (2007), 137–146.
8. Pariser, E. *The filter bubble: What the Internet is hiding from you*. Penguin UK, 2011.
9. Zhang, A. X., Ackerman, M. S., and Karger, D. R. Mailing lists: Why are they still here, what's wrong with them, and how can we fix them? In *CHI*, ACM (New York, NY, USA, 2015), 4009–4018.
10. Zhang, A. X., Blum, J., and Karger, D. Opportunities and challenges around a tool for social and public web activity tracking. In *CSCW*, ACM (2016).
11. Zhang, A. X., Culbertson, B., and Paritosh, P. Characterizing online discussion using coarse discourse sequences. In *ICWSM*, AAAI (2017).
12. Zhang, A. X., Igo, M., Facciotti, M., and Karger, D. Using student annotated hashtags and emojis to collect nuanced affective states. In *Learning @ Scale*, ACM (2017).
13. Zhang, A. X., Verou, L., and Karger, D. Wikum: Bridging discussion forums and wikis using recursive summarization. In *CSCW*, ACM (2017).
14. Zyto, S., Karger, D., Ackerman, M., and Mahajan, S. Successful classroom deployment of a social document annotation system. In *CHI*, ACM (2012), 1883–1892.