Building Systems to Improve Online Discussion

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

Abstract

Little has changed in the design of online discussion systems in the decades they have been available, even as problems involving scale, loss of context, and bad actors mount with broader use. To solve these problems, my research is on building novel online discussion systems that give users direct control over their experiences and information. Specifically, I focus on: 1) summarization tools to make sense of large discussions, 2) annotation tools to situate conversations in the context of what is being discussed, as well as 3) moderation tools to give users more fine-grained control over governance and delivery of messages.

Author Keywords

online discussion; forums; chat; social media; email; summarization; annotation; tagging; sensemaking; deliberation; moderation; online communities

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, anecdotes, humor, explanations, coordination, and deliberation [1]. 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
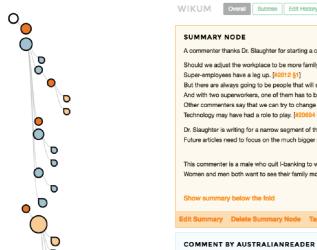
**Figure 1:** Wikum: a tool for summarizing long discussions.
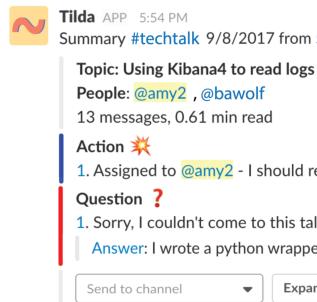
**Figure 2:** Tilda: a tool to mark up and summarize group chat.
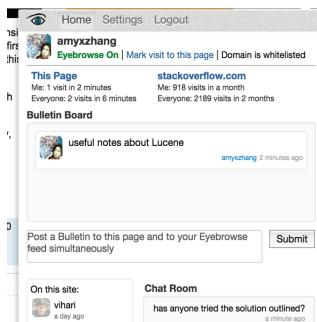
Figure 3: Eyebrowse: bringing social experiences to browsing.

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 discussion*, making it difficult for participants to get an overview of what happened or refer back. A related problem is the *loss of context* when discussions are stripped of information about their particular time, place, origination topic, or make-up of participants. Other times the problems are due to *bad actors* or *contested norms*, leading to tensions between participants and, in the worst cases, escalating into harassment of individuals.

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 and communities that give them greater control over their experiences and information online**.

Proposed Dissertation Work

My work can be divided into three main areas of inquiry that each examine design considerations around a particular technique or lever of control.

1) Summarization Tools to Make Sense of Large Discussions
The first set of systems help make sense of large discussions by providing collective **summarization** capabilities within discussion systems to provide a mechanism for people to deal with scale as well as a way to reflect on what was discussed. Because summaries are time-consuming and difficult to create, the project **Wikum**[1] [9] considers how a group of people could individually contribute small

---

[1]wikum.org

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 *recursive summarization*, 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. We are next researching Wikum's potential as a tool for Wikipedia editors to summarize discussions on Talk Pages [2]. A second tool **Tilda**[2] [6] focuses on summarization for group chat. The tool works by allowing participants to mark up an ongoing continuous chat stream to turn it into discrete structured summaries. This allows users to easily get an overview of the chat for when they fall behind. From interviews, we discovered what information users want to glean from chat as well as alternative designs for chat summaries, and then developed and deployed Tilda to several active Slack groups for work.

In relation to this work, I am interested in continuing to explore how humans and automatic techniques can work in tandem to construct summaries. My work on the characterization of different types of discussion using common discourse sequences [7] is one step towards understanding how summarization needs and workflows for creating summaries change depending on the type of conversation.

2) Annotation Tools to Situate Conversations in Context
The second set of systems situate conversations in the proper context by providing **annotation** capabilities when having discussions. One way to provide context with annotation is to allow discussions to live as annotations in the "margins" of other primary content, such as webpages, articles, or textbooks. An example is **Eyebrowse**[3] [5], a

---

[2]tildachat.com
[3]eyebrowse.csail.mit.edu
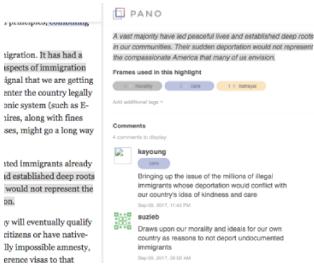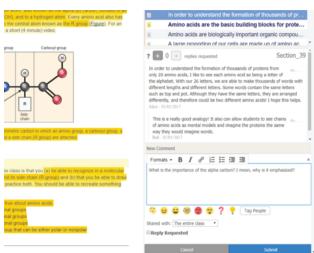
browser extension and web tool that allows users to share aspects of their browsing history, such as when they are active on particular sites. This allows people to "bump into each other" in spaces on the web and have serendipitous live conversations or asynchronous discussions that are then tied to the pages they are on. Another way to provide context with annotation is to annotate the comments themselves with relevant information. Here, I focus on self-annotation by commenters to both reveal the underlying basis of their comment and encourage self-reflection. For instance, **Pano**[4] is a tool that allows people to tag comments and articles with moral framing in order to become aware of their own and others' underlying moral values. A 10-day field study we conducted provided evidence that users improved their ability to frame arguments in the values of the opposing side.

Finally, systems like **Nota Bene (NB)**[5] [10] demonstrate that the combination of both types of annotation can be useful. NB allows 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 self-annotation of emotional state within a comment using emojis and hashtags could provide signals for readers about what parts of a textbook to focus on [8]. This could also help instructors cut through the noise of discussion to find unanswered questions or resolve disputes, or distinguish interesting from confusing passages for future revision.

3) Moderation Tools for Finer-Grained Control of Delivery
Finally, the third set of systems provide more powerful **moderation** tools that give users finer-grained control over the

Figure 4: Pano: annotation of comments with moral framing.

Figure 5: NB: adding emojis to comments to capture emotional reactions to textbook passages.

delivery of content, including what messages they see, who in turn sees their messages, as well as how messages get accessed and delivered. This can help manage information overload as well as help a community deal with contested norms. From studying existing communities, 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 [4].

One solution is to give users more fine-grained control over delivery of content, so that all members of a community are working together to ensure content gets delivered only to those who want to receive it. Drawing from our research into tensions within online groups, we built **Murmur**[6], 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 also study how novel moderation capabilities can combat problems such as harassment. We developed the tool **Squadbox**[7], where people facing harassment in their personal inboxes can recruit their friends as moderators [3]. From interviews with 18 recipients of online harassment, we learned how users wish to personalize their strategy for dealing with what they consider harassment, and designed Squadbox to be customizable to their needs.

---

[4]pano.csail.mit.edu
[5]nb.mit.edu

[6]murmur.csail.mit.edu
[7]squadbox.org
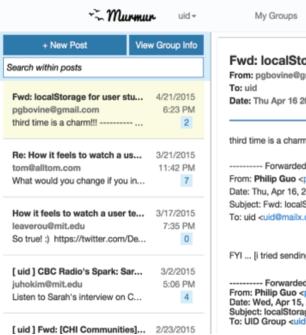
**Figure 6:** Murmur: configurable mailing list system.
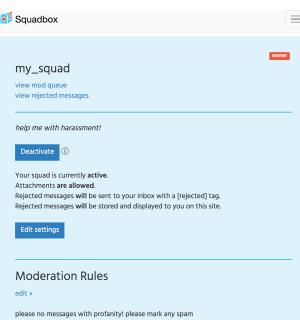
Figure 7: Squadbox: combating harassment using friendsourced moderation.

Expected Contributions

In my dissertation research, I hope to make the following contributions: 1) learn from specific communities of people engaging in online discussion to understand their needs, problems, and existing strategies; 2) explore the design space of possible user-led and community-led controls to enrich, organize, and govern discussion; and 3) produce design recommendations for social media platforms and communities to support better conversations at scale.

Goals for CSCW Doctoral Colloquium

One of my goals is to gather feedback on the direction and framing of my thesis topic as it stands. In my Ph.D. thus far, I have worked on a number of different projects and am now reflecting on both the coherence of my work and whether there are significant gaps in what I propose to cover. In the final year of my program, I am interested in suggestions on new work or extensions of existing work I should embark upon to narrow down or round out my current direction of inquiry. A second goal is to broaden my horizons to make interesting connections to academic disciplines, communities, and research problems that are adjacent to my work. I would like to improve how my work points to and sits in context to historical and existing lines of research.

REFERENCES

1. Justin Cranshaw and Aniket Kittur. 2011. The polymath project: lessons from a successful online collaboration in mathematics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1865–1874.

2. Jane Im, Amy X. Zhang, Christopher J. Schilling, and David Karger. 2018. Deliberation and Resolution on Wikipedia: A Case Study of Requests for Comments. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 74 (Nov. 2018), 24 pages.

3. Kaitlin Mahar, Amy X. Zhang, and David Karger. 2018. Squadbox: A Tool to Combat Email Harassment Using Friendsourced Moderation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 586, 13 pages.

4. Amy X. Zhang, Mark S. Ackerman, and David R. Karger. 2015. 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, 4009–4018.

5. Amy X. Zhang, Joshua Blum, and David Karger. 2016. Opportunities and Challenges Around a Tool for Social and Public Web Activity Tracking. In *CSCW*. ACM.

6. Amy X. Zhang and Justin Cranshaw. 2018. Making Sense of Group Chat through Collaborative Tagging and Summarization. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 196 (Nov. 2018), 27 pages.

7. Amy X. Zhang, Bryan Culbertson, and Praveen Paritosh. 2017a. Characterizing Online Discussion Using Coarse Discourse Sequences. In *ICWSM*. AAAI.

8. Amy X. Zhang, Michele Igo, Marc Facciotti, and David Karger. 2017b. Using Student Annotated Hashtags and Emojis to Collect Nuanced Affective States. In *Learning @ Scale*. ACM.

9. Amy X. Zhang, Lea Verou, and David Karger. 2017c. Wikum: Bridging Discussion Forums and Wikis Using Recursive Summarization. In *CSCW*. ACM.

10. Sacha Zyto, David Karger, Mark Ackerman, and Sanjoy Mahajan. 2012. Successful classroom deployment of a social document annotation system. In *CHI*. ACM, 1883–1892.