tldr

INTERFACES FOR LARGE-SCALE ONLINE DISCUSSION SPACES

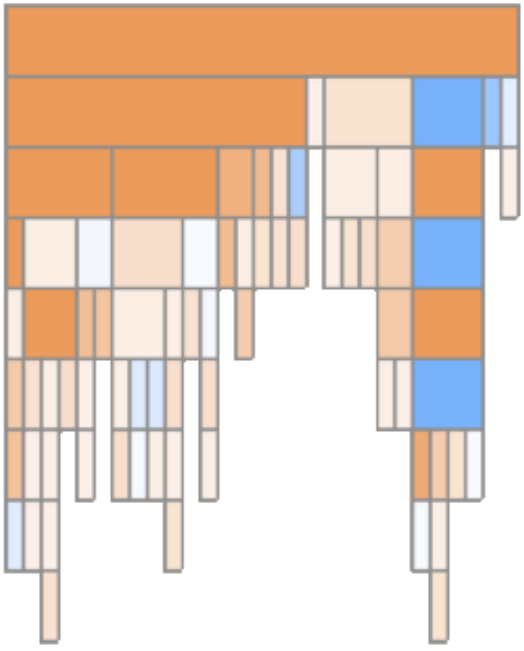

MIMS Final Project by
Srikanth Narayan

under the guidance of
Coye Cheshire

Abstract

Recent years have seen a proliferation of online discussion spaces with hundreds to thousands of messages/participants. Users participating in these discussions are overwhelmed by the sheer amount of information presented, and the systems that support these conversations are lacking in functionality that lets users navigate to content of interest. This paper discusses the design of an experimental interface for navigating through large-scale discussions. Guided by user research and data analysis, the design incorporates several unique features such as visual overviews, appropriate social cues, non-linear navigation and multidimensional filtering, all of which help users find information that is of most interest to them. Preliminary user testing results suggest strong interest in using the application for participating in discussion spaces.

Introduction

Online discussion spaces are used for shared public interpersonal conversations between groups of people. These discussion spaces are an asynchronous communication medium, and thus enable interaction between participants without requiring continual presence. Most discussions are user-initiated, user-maintained and user-moderated. A user can essentially reach a wide audience through this system, as well as listen to messages from this range of users. They foster discussions on every imaginable topic, and because of its wide reach, bring in opinions from various perspectives. In essence, they help in creating a sense of communion between globally separated groups.

From the days of the USENET and Bulletin Board Systems to the present day, discussion spaces have seen widespread adoption on the internet. In an effort to encourage active user participation, more and more web destinations have started to incorporate a discussion component into their system. Discussions on popular websites grow proportionally with traffic, and as a result, it is not uncommon to encounter discussions with hundreds to thousands of messages/participants. This phenomenon can be observed on a wide variety of websites - news outlets, blogs, social media websites, community websites and support forums. While most of these discussion spaces are able to support small discussions, their effectiveness is greatly reduced as the discussions grow larger. Communication systems that have supported interactions in such large numbers are relatively novel, and the dynamics of such conversations are poorly understood. Even from a cursory examination of interfaces presented by these systems, one can notice that they vary greatly in how they support these conversations - something that hints at an overall ill-understanding of the problem.

Though discussion spaces have grown to become a heavily-used communication medium on the internet, it isn't a technology that has seen

continual cycles of improvement. Even today, most websites present discussions in an unusable linear list of messages. phpBB continues to be the most popular option for web based forums despite its ineffectiveness in handling discussions. As these discussions grow larger, the limitations of the interfaces for navigating these discussions exacerbate. Users are overwhelmed by the sheer amount of information presented, and without an effective way to filter through the information, tend to become overloaded. This paper describes a project aimed at investigating effective designs for navigating through large-scale discussion spaces. The paper starts out by contextualizing the problem of large-scale discussions, and explicates on how it could be lead to a disruptive user experience if not handled appropriately. Following that, various stages of the design process are elaborated. Results from a preliminary round of user testing is presented in the end.

Related Work

The topic of large-scale discussions has been one that has interested researchers since the emergence of USENET, and there have been several research efforts aimed at understanding the inherent dynamics of such conversations. Some projects have been analytic in nature, while others have sought to design better end user interfaces. In this section, a selection of related previous work, along with key findings is discussed.

Sack's Conversation Map[1] provides a visual representation of threads in USENET over time, combining it with social network of interactions within the conversation, and a semantic network of themes within the conversation. It also served as a tool for exploring the conversation, allowing the user to bring up the content of individual messages. The visualizations were presented as the primary interface, and the content of the discussion was always a layer beneath the visualizations.

Loom[2], a project done by Donath et al. visualizes online discussions on USENET, with the primary goal of uncovering patterns within these discussions. Temporal distribution of a participant's contribution, patterns within sequence of messages, and patterns based on content of messages were some of the results of this project.

Smith et al.[3] detail another system that highlights structural and temporal patterns within USENET conversations using thread tree visualization, along with a visualization of the social network within the conversation. Overall goal of the project was to develop visualization components for threaded conversations that enable deeper understanding of patterns of activity.

Another project by Dave et al.[4] is one more closely related to the project described in this paper. The focus of their research was large-scale, topic-centered, transient discussions that they termed as flash forum. Using data

from the popular news aggregator Slashdot as an example, they detail a system that combines data visualization and text analytics in a novel interface for large-scale discussions. This work sheds light on several nuanced distinctions between traditional discussion spaces such as USENET when compared to flash forums such as Slashdot. An effective visualization for discussion spaces is introduced. Comparison studies of text analytics and visualization are presented as well.

In summary, most of the research efforts have centered around building analytic interfaces/visualizations, rather than interfaces that could be deployed to the end-user. This project differs significantly from many of the earlier attempts in that respect - the goal of this project is to propose designs that would work well for users participating in discussion spaces. Another important distinction to be noted - majority of the previous research on large-scale discussions has been done on data from USENET/mailing list conversations. Discussions afforded by newer social media websites such as Reddit are considerably different from the conversations on USENET, and are more closely related to flash forums[4]. As a result of the advanced functionality provided by Reddit when compared to other discussion spaces used in previous research, the data that can be extracted out of discussions on Reddit is a lot richer. Unambiguous hierarchy of messages, rich two-way moderation data, author history, story categorization and high precision timing are some of the unique features of the dataset extracted from Reddit. Based on this fact alone, there is an opportunity to uncover more meaningful patterns within conversations, and possibly understand the dynamics of these conversations better.

Understanding Large-Scale Discussions

In addition to being heavy traffic websites, there are other characteristics of these discussion spaces that foster growth to an extent that is of interest in this project:

- Translucency: Activity in the discussion space is discernible by everyone, i.e., many of these spaces maintain a level of translucence to the outside world that is vital in attracting more people into participating in them.
- Low barrier to entry: Upon finding something interesting that a user would want to contribute to, it is relatively easy to gain access and start posting to the discussion space.
- Anonymity: While most of our day-to-day interactions are with people we are acquainted with, discussion spaces are unique in the way that most interactions are between complete strangers. One's own identity does not matter, and is usually known only by a consistent pseudonym.
- Transience of discussions: These discussion spaces are marked by a high churn of discussions that surface over time. This is influential in inciting users to keep returning to the discussion space.
- Persistence of messages: Being an asynchronous form of communication, messages should persist long enough so that the participants in the discussion get to read it. In most discussion spaces, messages persist

seemingly forever, growing quickly into a repository of discussions that users can peruse and reinstate.

As these discussion spaces grow larger, common problems that manifest themselves include inability to find messages of interest, low signal-to-noise ratio, high redundancy, missed communication inputs and many more. Though the amount of information streaming through a discussion space is large, lack of a way to control and filter through the information exacerbates the problem even more [7]. In their work on structuring CMC systems to avoid information overload[9], Hiltz et al. declare that "The value of an information system lies in what it withholds, as much as in what it gives". As human beings, we are governed by basic limits in our abilities to process information. As these discussions grow larger, the amount of information that merits attention can exceed the ability of the user's ability to process it[12]. The sheer amount of information could potentially overwhelm the participants, leading to individuals becoming overloaded. A natural instinct of human beings is to resort to the path of least effort[8], and users are more likely to develop strategies to evade being overloaded[10]. Previous research that has studied the effects of Information Overload on participants in a discussion space has found that it leads to decreased participation, or even an end of active participation[11].

First introduced by Hiltz & Turoff[9], the term Information Entropy refers to the lack of sufficient organization of messages to be recognized as significant or as part of a conversation's history. This lack of distinction can be attributed to be one of the primary causes of the perceptions of information overload, and long been speculated by researchers as a opening to attack the problem of information overload. This premise guides the solution presented in this paper as well.

Goals of this project

The previous section briefly explicated how inefficiencies in handling large-scale discussions could potentially lead to gradual breakdown of a discussion space. As stated earlier, the main goal behind this project was to investigate designs that could combat the aforementioned problems successfully.

To achieve these goals, a thorough understanding of milieu of discussion spaces is essential. Reddit, a social news aggregator website that regularly hosts large conversations, was chosen as a model for understanding the inherent dynamics in these conversations, both for user research and data analysis. Design options generated from these findings were to be tested with end users, and through iteration, arrive at design solutions that would be successful at handling these discussions.

A Note on Reddit

Reddit [at reddit.com] is a popular social news aggregator web application. Users of this website submit stories of interest to the site, and stories that are popular among its users are promoted to the frontpage, a part of the website accessed by most of its users. Each story submission has an associated 'Comments' component where users hold discussions about the topic of the story. As a result of high activity on the frontpage, it is not uncommon for these discussions to reach a few hundred to thousands of comments.

A brief discussion of the features offered by Reddit's discussion interface would be beneficial in the context of discussion of this project. Reddit's discussion interface is one that has been hailed by its users as one of the best on the internet. It offers features such as message threading, two-way moderation, and message sorting. The interface is extremely simple, primarily textual and content-rich. Filters such as message limit, moderation score limit and thread collapsing help the user control the number of messages displayed.

Despite being a highly usable interface, the limitations of the current Reddit interface as the discussions grow larger are apparent. As the comment threads are displayed in a linear expanded list, interacting with the interface requires an inordinate amount of scrolling. For example, a discussion capped at 500 messages contains enough information to fit 60 fullscreens on a 1024x768 display i.e. the user would have to scroll through 60 screens of text to read through the discussion. Displaying the entire list of comments at once also results in extremely slow page loads and page scrolls. Reddit runs a hotness algorithm on comments in the discussion space, and the arrangement of threads is often the most relevant information guiding the user in navigating to messages of interest. Limitations of this interface are

applicable to other websites such as Digg/Slashdot etc that support similar conversations as well.

User Research

To gain a better understanding of users' perceptions and behavior on discussion spaces, a survey instrument was designed. The survey sought to collect information about an individual's level of participation, activities, interests and preferences with respect to discussions spaces. The survey was advertised with help from the founders of Reddit, who generously ran a frontpage ad linking to the survey. Over a period of 15 days that the survey was active, there were 477 responses from users of the website. Responses came in from all over the world, and most respondents were enthusiastic about offering thoughts on their usage of discussion spaces. About 55% of the respondents chose to answer the optional question on their thoughts on how discussion spaces could be improved, with some leaving detailed multi-page responses on problems they encounter. About 25% of the respondents chose to leave their personal contact information out of interest in the project. Both of these numbers hint at the high quality of data collected from the survey.

In this section, some of the key findings from the survey, and their implications for design are discussed:

- A majority of the respondents (68%) have been users of Reddit for more than 12 months. This gives us an indication that most of the users on the site have been there for long, and that the site has a steady and loyal userbase.
- With regard to activities that the users are involved in on the site, the survey asked the respondents to estimate their extent of various activities within a discussion space.
    - More users are involved in responding to comment threads rather than starting one. 37% of the respondents indicated that they rarely start comment threads, whereas 42% of the respondents indicated that they are occasionally involved in responding to comment threads.
    - Most of the users are never involved in moderation of comments. 52% of the respondents indicated that they never moderate or are rarely involved in doing so. This implies that all the moderation of the website is done by half of the userbase, and the other half is rarely involved in this activity.
    - A majority of the respondents (92%) indicated that they are often involved in reading comments (lurking). This indicates that this is the predominant activity of users in a discussion space, and an interface built for discussion spaces should be optimized for supporting this activity.

- Overall summarization of the responses is shown below

Starting Comment Threads	15.9% 76	39.2% 187	31.2% 149	11.3% 54	2.3% 11
Responding to Comment Threads	4.6% 22	16.6% 79	41.3% 196	24.8% 118	12.6% 60
Moderating Comments	42.0% 200	10.1% 48	16.8% 80	17.9% 85	13.2% 63
Reading Comments (Lurking)	0.2% 1	1.3% 6	7.1% 34	19.5% 93	71.8% 342

- With regard to their extent of involvement in the discussions on Reddit, 50% of users indicated that they participate in discussions regularly (More than 3-5 days a week).
- With regard to qualities in comments that are of interest to them, comments that are insightful are the ones that are the most preferred, followed by Informativeness and Humor. Respondents also added in other adjectives such as Intelligent, Corrective, Trolling, Puns, Relevant and Well-written.
- In trying to assess what factors influence their decision to read a particular comment, the survey question was designed to assess several factors that could potentially be influential. The results proved to be surprising, and quite eye-opening:
    - Moderation score was found to be a good estimator of worthiness of a comment, with 65% of the respondents indicating that they found it to be important to some degree.
    - Comment Author was found to be the least effective at indicating quality of a comment. Only 17% of the respondents thought that Comment Author was of any importance at all.
    - Length of the Comment, and Number of Replies that a comment has received are equally good indicators of whether a comment is going to be read. About 50% of the respondents indicate that it is important to some degree.
    - Position of the message in the discussion space was found to be of high importance with the respondents, with 75% of the respondents indicating that it is important to some degree. This denotes that most users rely heavily on the reddit hotness algorithm that pushes the best comments to the top.
    - Predictably, the best indicator of worthiness of a comment was found to be the user's perception of the message after reading the

first few sentences. This corroborates an observation made during the user interviews. Most users are skimming through messages in a discussion space, going through them at a rate faster than normal reading. Only when they find something that is of interest, do they slow down to read it more thoroughly.

- Overall summarization of the responses is shown below

Moderation Score	7.4% 35	14.5% 69	13.5% 64	47.6% 226	17.1% 81
Comment Author	48.9% 232	15.8% 75	18.4% 87	13.5% 64	3.4% 16
Length of the Comment	4.7% 22	16.7% 79	28.3% 134	43.8% 207	6.6% 31
Number of replies	7.4% 35	20.7% 98	24.1% 114	37.6% 178	10.1% 48
Position in the discussion space	4.2% 20	7.2% 34	13.1% 62	38.7% 183	36.8% 174
Perception of the message after skimming the first few sentences	2.3% 11	4.4% 21	8.7% 41	43.6% 206	41.0% 194

- Another surprising result from the survey was that 60% of the respondents indicated that they do not use filters to limit the number of comments displayed in the discussion space, but rather try to find interesting comments within the entirety of the discussion. This is probably due to the design of the Reddit system which forces the user to make system-wide settings rather than applying them dynamically within the discussion space. Among users who do use the filters, Moderation score limit was found to be the most favored. Ability to collapse entire message threads and Message limit are the other two filters supported on Reddit, and they were preferred by the respondents in that order.

Data Analysis

In addition to user research, the design was also informed by analysis of the data within large-scale discussions. While there has been much research that has focussed on insight that could be gained from mining corpora of large number of discussions, the efforts for this project were mainly focussed on understanding data that comes from within a *single* large-scale discussion.

Visualizations that could aid in *exploratory analysis* of data within discussions were developed, rather than relying on statistical means to analyze the data.

The first of the visualizations was built to corroborate an observation that was first made by Dave et al[4] with respect to flash forums. In their definition, flash forums are typically characterized by a short lifespan of intense activity. This holds good for data from Reddit as well, as can be seen from the visualizations shown below. In these visualizations, each vertical represents an hour, and all the comments made within that hour fall into that particular vertical. From the visualizations, it is easy to see that most of the discussions on the site typically last only for a few hours. This period is often marked by intense activity, which results in the discussion growing to hundreds of comments in a very short duration. In the visualization, the nodes are colored by their moderation score - with orange indicating hotness/positive moderation, and blue indicating coldness/negative moderation. This makes it easy to notice that many of the earlier comments fare well in terms of moderation score, while most of the later comments do not perform as well. A gradual fading of color over time indicates that it is often hard for a comment worthy of merit to be recognized as such, unless it is also posted at the right time.

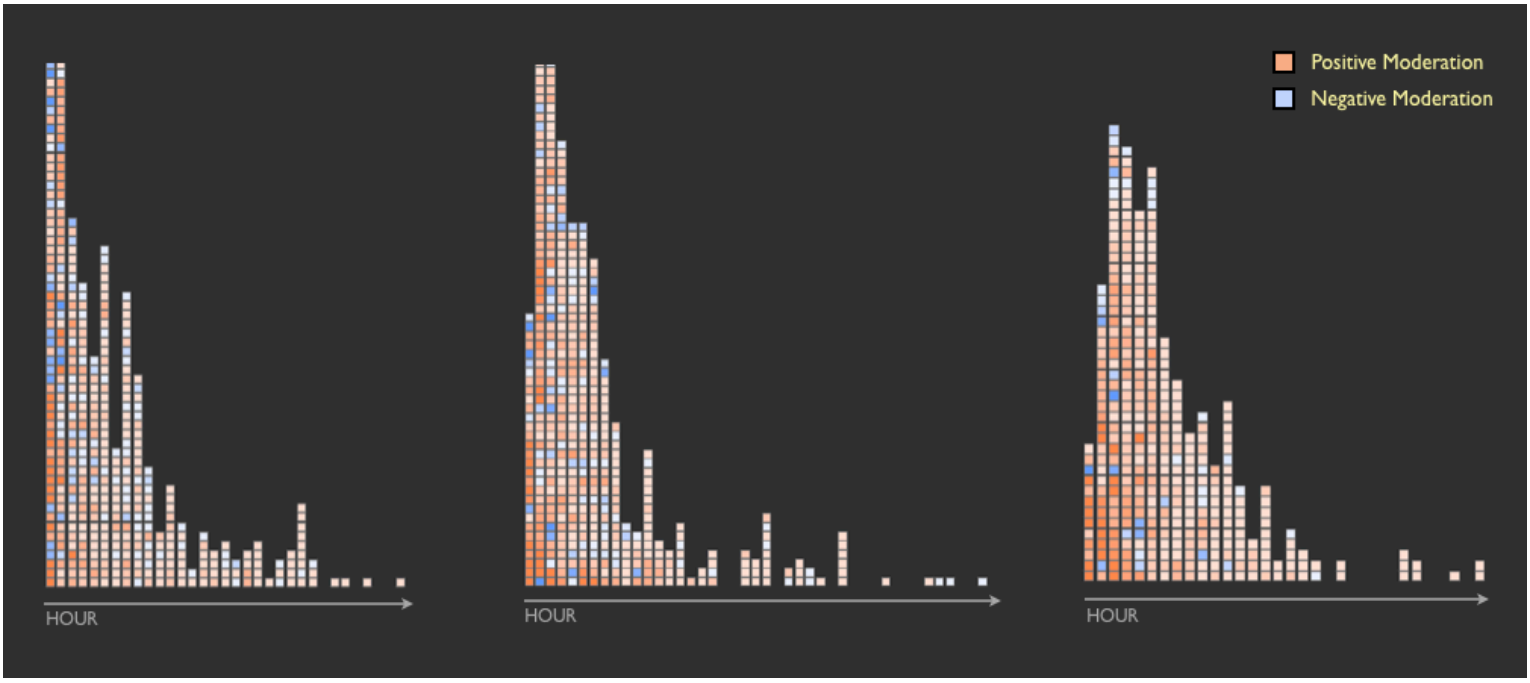

The next visualization was aimed at uncovering patterns within conversational exchange in a discussion. While it would be an arduous task to estimate the quality/nature of a message by analyzing the content, patterns in moderation score were found to be rough indicators of these two dimensions in certain cases. The visualization represents a discussion as a tree, with each node representing a comment in the discussion. Upward-facing triangles shaded with orange indicate positive moderation, and downward-facing triangles shaded with blue indicate negative moderation. Several patterns can be inferred from an examination of the visualization shown below. First of all, thread (2) seems to have a broad ongoing discussion. Certain branches of the tree tend to be more favored by readers than other branches (a), and that certain branches of the tree grow deeper than other branches (b).

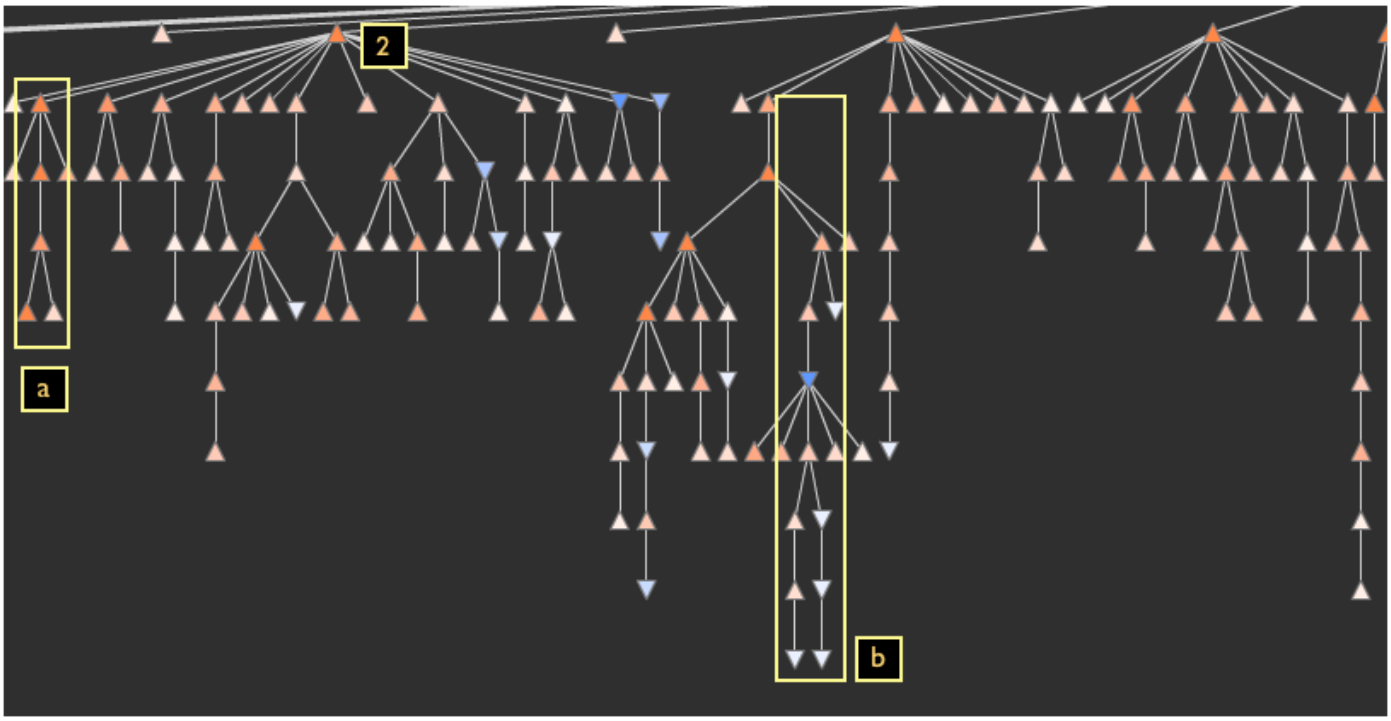

Examining another visualization shown below, other patterns within the data are uncovered. Alternating orange and blue nodes (b) could potentially hint at an ongoing argument, with most of the moderators siding towards one side of the argument. Continuous lines of blue(a) could potentially indicate trolling/unworthy messages.

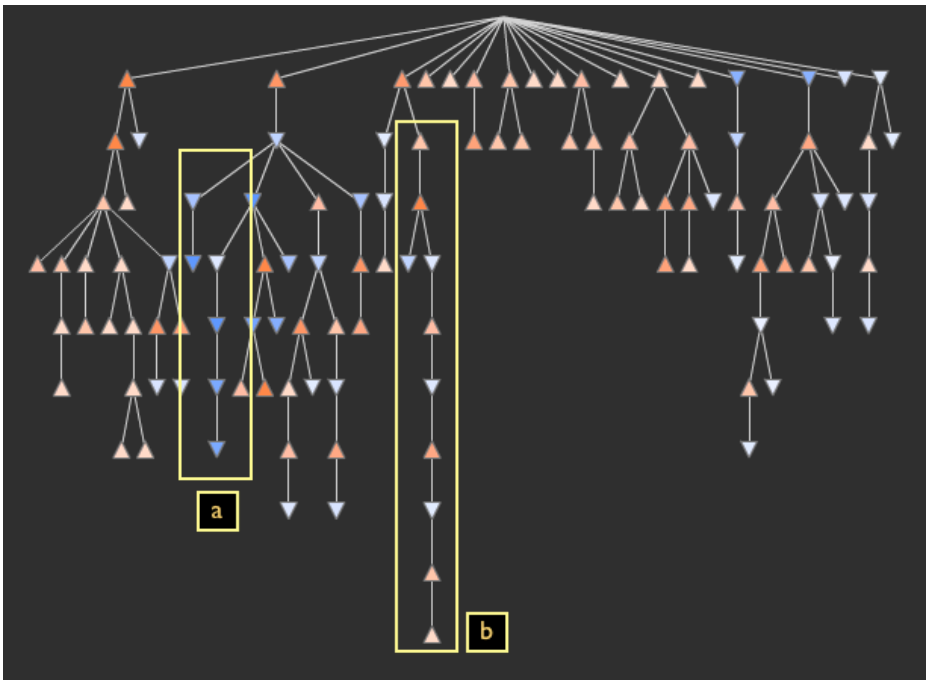

The patterns shown about could be made conspicuous by transforming the tree to a visualization shown below. This visualization, based on icicle tree layout, has nodes of unit length at the leaf level, and each parent has a width that is the sum of the width of its children. Each individual thread is thus represented by a first-level block, and all the messages within that thread are "fitted" below that block. Colors retain the same meaning as they did in the previous visualization.

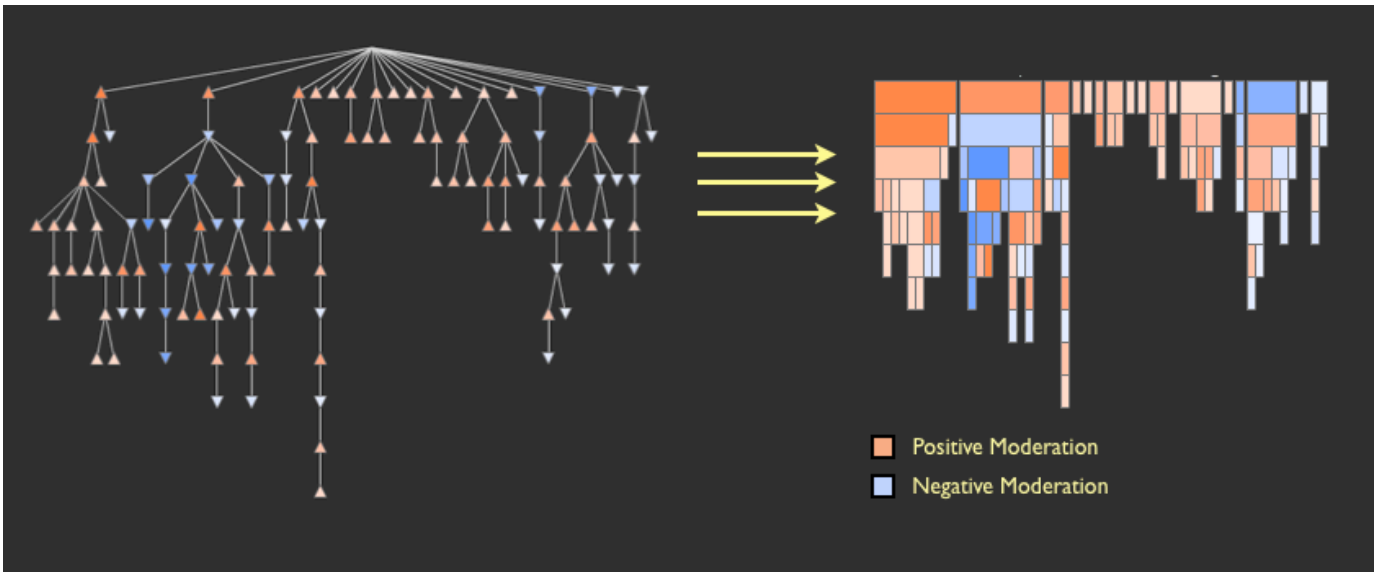

The next visualization was designed to facilitate comparison across threads. Activity, moderation trends, length of comments, and level of reply are important variables that help in comparison of threads. The image below shows a visualization that is an abstract representation of the indented comment thread seen in most discussion spaces. This visualization, which is a slightly improved version of a visualization proposed by Wattenberg[13], brings out several patterns within discussions. Firstly, it makes it easy to compare length and activity across several threads, as well as length and activity of comments within a single thread. We can see from the visualization that the first two threads (a & b) are much longer than most of the other threads. It is also easy to discern threads that have verbose comments(b) as well as threads that have short/terse comments(a). Progression of activity across threads can also be seen, which hints at how the placement of threads within the discussion space has affected its activity rate.

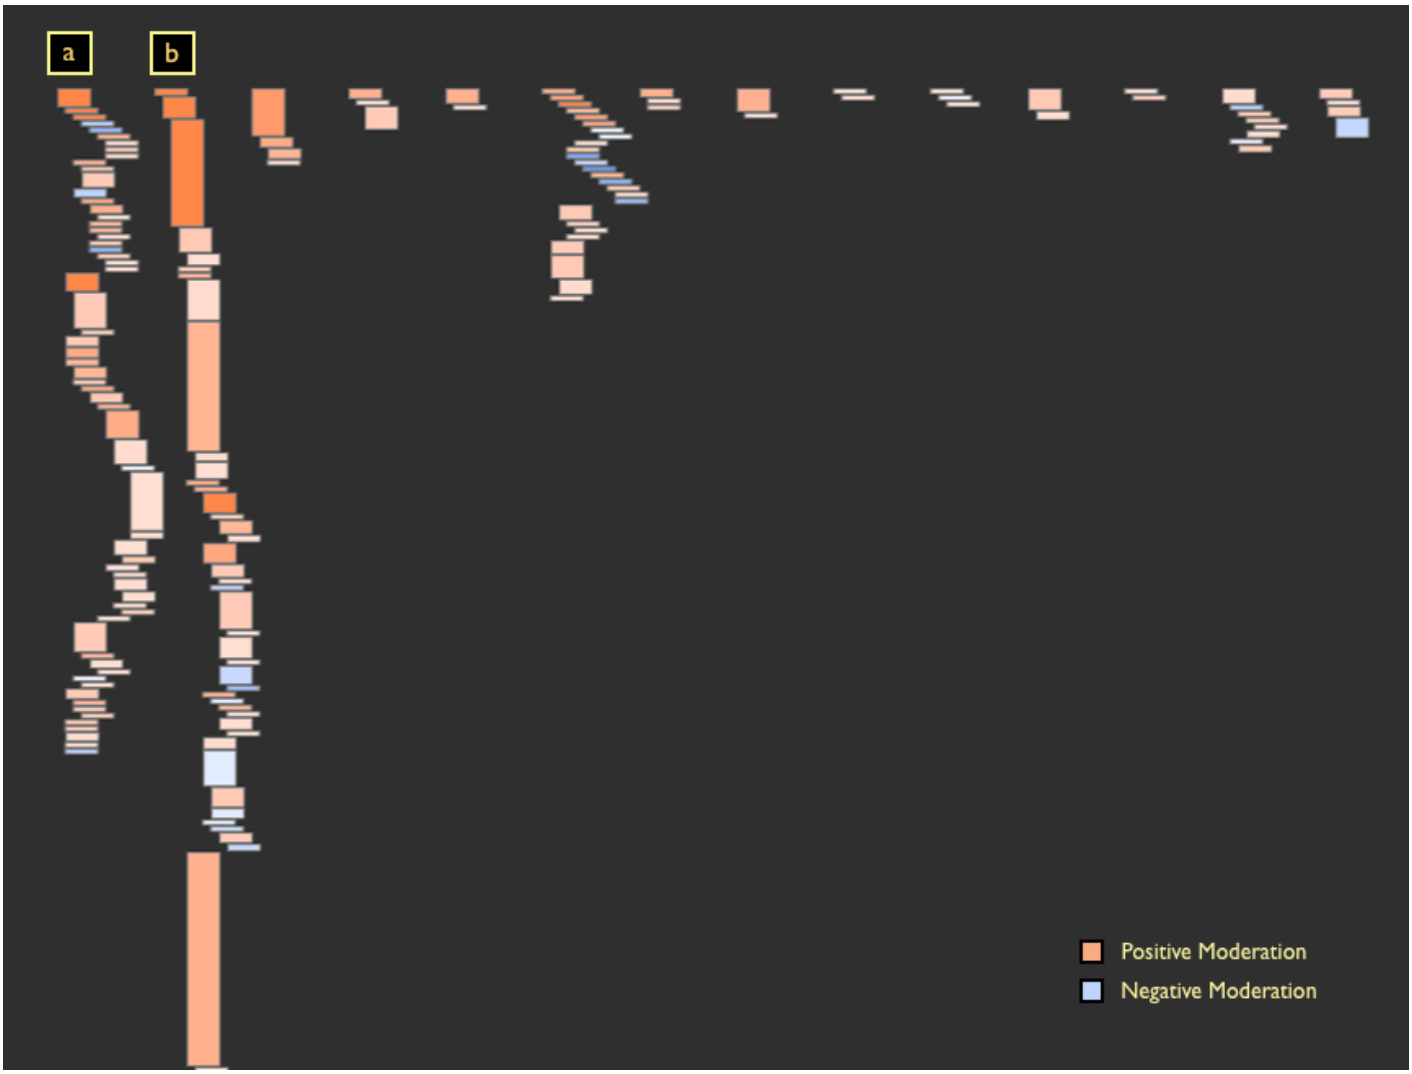

Design

The design of the application benefited greatly from the dual-pronged approach that was taken for needs assessment. As discussed earlier, needs assessment was not limited to traditional user research methods, but also incorporated analysis and understanding of data within large-scale discussions. Inspiration for design ideas came from both of these methods, and this particular method choice proved to be quite fruitful in the end. Effective utilization of spatial, semantic and social navigation [5] were considered in developing design solutions. Design iterations were mainly done as quick experiments, producing sketches, mockups or prototypes of ideas, and testing them for feasibility. In this section, the design behind the interface will be explicated in detail. Though the design went through several iterations, the discussion that follows is mainly centered around the last iteration of the design.

The attached image illustrates the primary interface for the discussion space. In this view, the user is able to get a sense of ongoing activity within the discussion space, and dive into parts of the conversation that interests him the most. Several elements presented in the interface aid the user in gaining an understanding of the ongoing discussion:

- Each block in this view represents an ongoing thread of conversation. Information that gives an overview of ongoing activity within a thread is presented in this view. The initiating message of a thread, which plays an important role in shaping the conversation, is presented at the top. Qualitative tags describing messages within the thread, a list of prolific authors, author count and reply count are other variables that are presented. Based on information provided in this view, a user can make assessments about their level of interest in the conversation.
- Along with this overview, a visualization that indicates how the conversation develops from the initiating message is presented for each ongoing thread. This visualization is based on a visualization discussed earlier in the paper, and reveals patterns in activity, message length and moderation. By presenting it in the context of each thread, the visualization supplements the information that is provided in the thread overview, and helps in making a better assessment of worthiness of a conversation thread.
- Upon finding a thread that the user is interested in, the user can dive into messages within that thread by 'expanding' the overview. The overview is replaced with an expanded list of messages presented in an indented tree view fashion that has become a standard view for representing message threads. Typography and structure have been optimized in this view to enhance readability of messages as shown below.

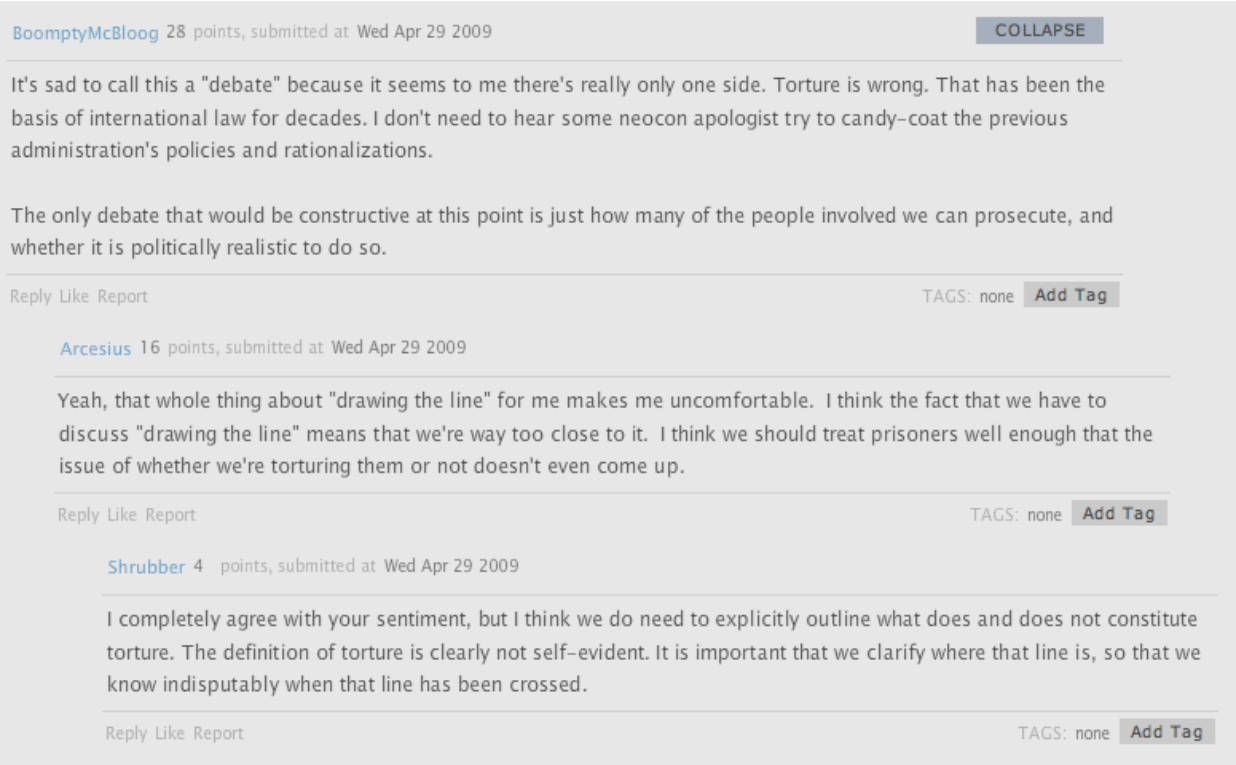

As seen earlier, the interface also provides an option to open a 'visual overview'. Upon opening this view, the interface seen in the image below is presented. This view presents several visualizations that depict activity within the entire discussion. These visualizations not only make patterns within conversations obvious, but also serve as a navigational aid to dive into particular points of interest. This view also presents several filters that enable users to limit the messages along certain dimensions. Overall, these set of functions were deemed as auxiliary to the core functionality provided by the application, and a design decision was made to hide the overview panel by default. When the user wants an overall picture of the discussion to navigate to points of interest, they can invoke this view.

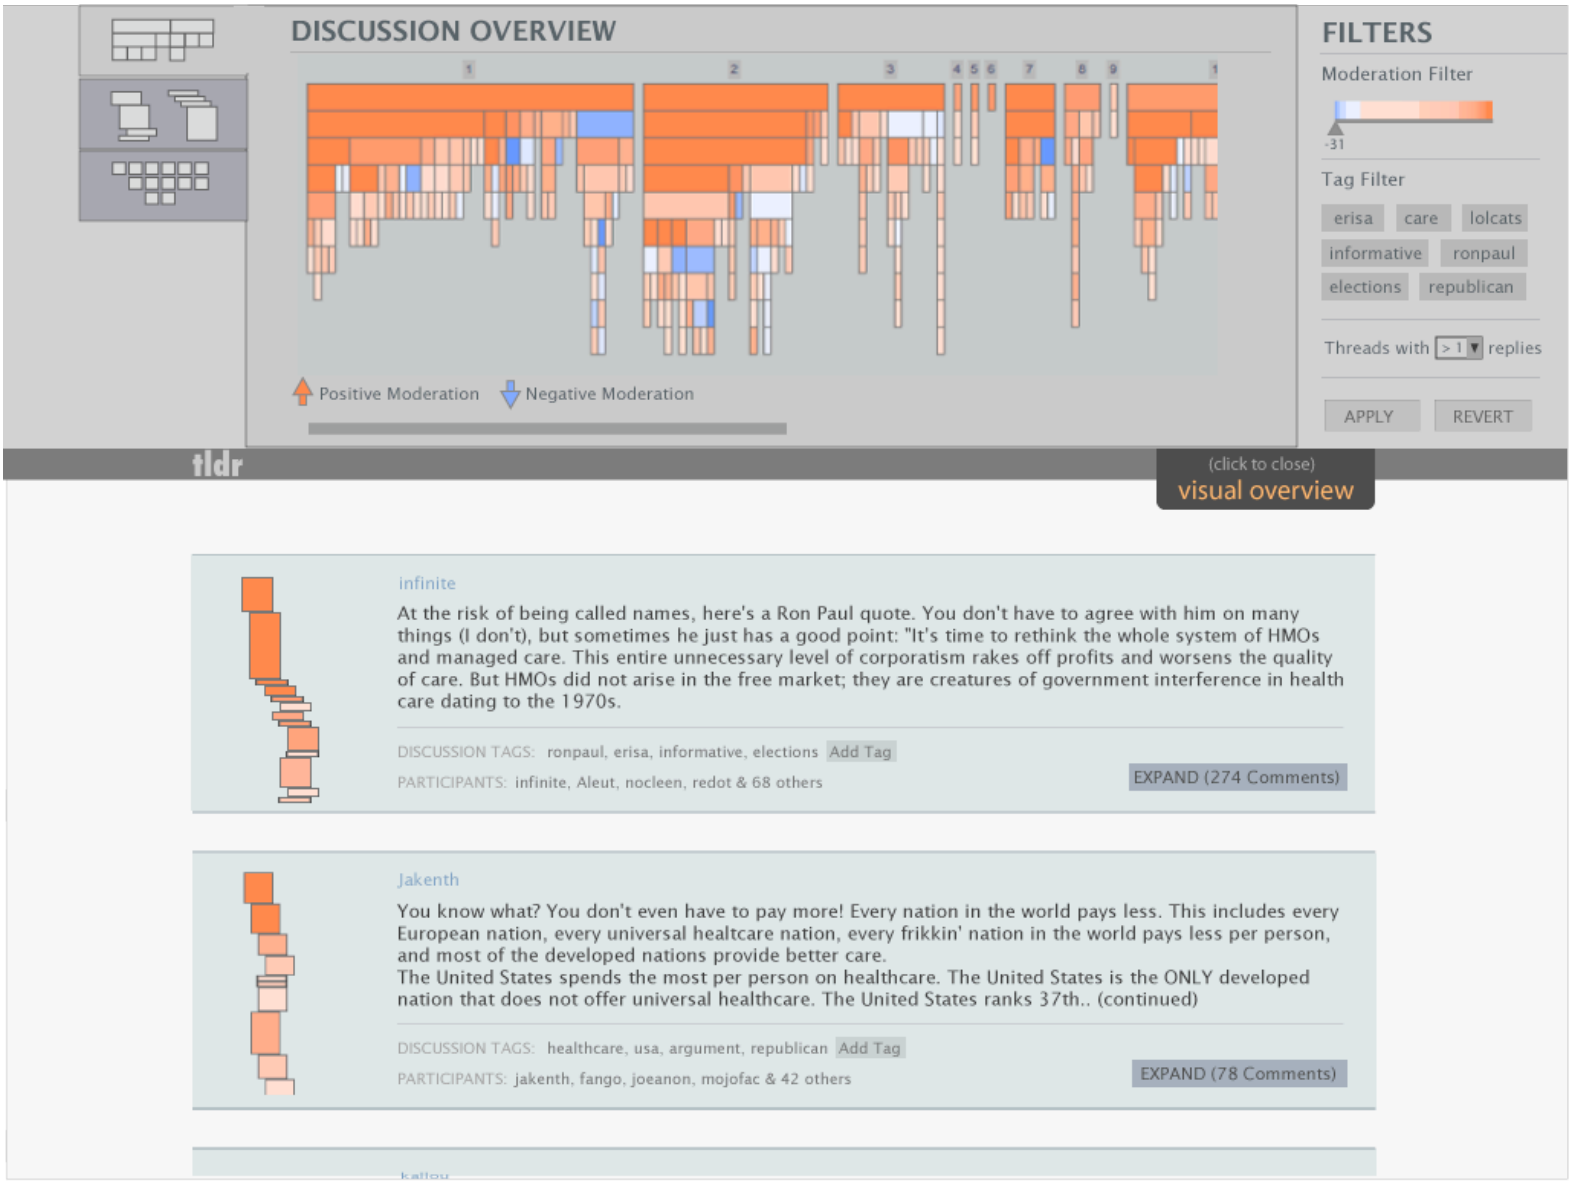

A detailed discussion of the design of elements within this view follows:

- Discussion Overview: Visualizations that represent activity within the entire discussion is presented in this view. The initial visualization is an adaptation of a visualization introduced earlier in the paper, and reveals several patterns within the discussion. It is easy to spot threads with the highest participation, threads that run the deepest, conversations within threads that have been favored the most, controversial areas of the conversation, and overall distribution of activity across threads. Upon hovering over any node within the visualization, relevant information about the particular comment is presented. Clicking on any node takes the user directly to the particular comment in the conversation view, presenting the comment in the context of the conversation. In this way, the visualization serves as an Overview + Detail interface. Another visualization within the Discussion Overview presents a comparison of threads similar to one shown earlier in the Thread Comparison visualization. This visualization serves as a navigational aid as well, allowing users to jump directly to messages of interest.
- In addition to these visualizations, this view also presents the user with several filters to limit the messages presented along a dimension of interest:
    - Moderation Filter: This filter allows the user to limit the comments displayed based on the moderation score it has received. User interacts with a simple slider, and receives immediate feedback through an updated visualization. A visualization placed above the slider presents the distribution of moderation scores within the discussion, and helps the user in an assessment of comments to filter out.
    - Tag Filter: This filter allows the user to select comments within a discussion that have been tagged with a particular qualitative descriptor. Every selection highlights messages with that particular tag in the visualization, and applying these settings updates the text view to open up the conversations that have been highlighted.
    - Reply Number Filter: The last of the filters allows the user to limit the thread based on number of replies that the thread has received.

Prototype Implementation

The prototype was implemented using Adobe Flex and the Flare Visualization toolkit[14]. This choice of technologies enabled development of an intricate piece of software in a short development time of 2 months. Though there was significant development to be done for building the visualizations, Flare provides an exemplary framework to build these visualizations upon.

As mentioned earlier, a decision was made to use the data from discussions on Reddit for the purposes of the project. Reddit provides a JSON API endpoint for its 'Comments' component. This JSON data provides rich information about the story topic and its subsequent discussion. Using a recursive module, it is possible to retrieve the discussion for any story in its entirety. Upon retrieval, this data is processed on the client, and an alternate representation based on the designs described above is rendered.

All the visualizations employed in the application have been developed as Flare layout classes. Further, the threaded conversation view (textual representation of the discussion) has been implemented as a Flare visualization layout as well, so as to enable fluid interaction between the visualizations and the textual representation. Data management within the application has been carefully optimized, which enables quick processing and representation of data structures with thousands of nodes. As a result of these techniques, the response from the application upon user action is quick, even though some actions are fairly computationally intensive.

In its final version, the application is built to be deployable on the web. Any user of Reddit who is interested in browsing through the discussion using this application can do so easily. This also enables the ideas developed in this application to be tested for effectiveness across a wider user base.

Evaluation

To evaluate the efficacy of the application in aiding navigation through large-scale discussions, a user study was conducted. This should not insinuate that the user feedback was sought only after the development of the application. Strict user-centered design process was followed, and user feedback was instrumental in guiding design decisions that led to the current implementation.

In the initial round of user testing, six participants were studied as they used the application. Though none of the study participants were users of Reddit themselves, they were all users of fairly large discussion spaces such as Slashdot, deviantART, vBulletin etc. Most of these participants were heavy users of discussion spaces, with five participants indicating that they use one everyday. They have also used these discussion spaces for longer than a year.

The study itself was structured with two goals in mind. To study ease of understanding of the interface for first-time users, and to test the effectiveness of the interface in achieving specific tasks. Each study was conducted over a period of 30-45 minutes. Initially, the users were asked to explore the functionality offered by the application, and a think-aloud protocol was employed to gain an understanding of how users perceived the functionality they were presented with. Following that, participants were asked to perform several scenario-based tasks that tested their understanding of specific features. For example, one of the tasks tested their ability to identify the best messages in a large thread using the visualization. Another task tested their ability to identify threads with lengthier postings.

Overall, the user testing proved to be extremely fruitful in uncovering usability bugs in the application. Jakob Nielsen has pointed out that usability testing with 5 users uncovers most usability problems with the application[15], and the number employed in this study was found to uncover several unique usability problems. While some of the problems found were a result of incomplete implementation of the application, the test was useful in uncovering overseen usability issues, and helping refine the design for easier understanding. Majority of the participants were able to gradually understand, through exploration, all the novel visualizations and functionality presented in the application without any guidance from the moderator. Even when the participants themselves sought guidance, hints were provided to guide them, rather than providing them with particular answers.

All of the participants had positive feedback about the application. Several of the participants used the word 'Awesome' as they were exploring the application. Some of the other descriptors used during the think-aloud/post-study conversation included 'Coool', 'Very Impressive', 'Most meaningful visualizations I've seen (for discussion spaces)' and 'State-of-the-art'. Most of the participants were also fairly interested in using the application, and some interacted with the tool beyond the requirements of the task. Among the visualizations presented, several participants mentioned that they found the icicle visualization more compelling, and that it made it easier to recognize patterns within conversations.

At the end of the user testing session, participants were asked to fill out an exit survey. Results from this survey are discussed in this section:

- First of the feedback questions asked them to rate the effectiveness of the features they had interacted with, on a five-point scale from 'Not at All Effective' (value=1) to 'Extremely Effective' (value=5).
    - Thread Overview feature, which presents a summarization of activity within each thread, received a score of 3.5 on a five-point scale.
    - Thread Thumbnail Visualization, which presents a visual overview of the structure of the thread, received a score of 3.33 on a five-point scale. This was the lowest rated feature of the application.

One user commented that the visualization required scrolling, an activity that he would want to avoid as much as possible.
    - Discussion Overview Visualization, which brings up a visual overview of the entire discussion, received a score of 4.33 on a five-point scale. Most of the users commented positively upon exploring the feature for the first-time.
    - Navigation using Discussion Visualization, which lets users navigate to certain points of interest, received a score of 4.17 on a five-point scale. Two of the participants noted that the context-highlighting and message folding were particularly useful.
    - Moderation Filter, which lets users filter the comments based on comment score, and receive immediate visual feedback, received a score of 3.5 on a five-point scale.
    - Tagging/Tag Filter, which lets users navigate to certain messages based on the qualitative tags assigned by other users, received a score of 3.5 on a five-point scale.
- The next feedback question asked them to rate the product along several dimensions. Again, the users were asked to rate it on a five-point scale from 'Strongly Disagree' (value=1) to 'Strongly Agree' (value=5).
    - The first question which asked them to rate the usefulness of the application, received a score of 4.5 on a five-point scale.
    - Ease-of-use of the application received a score of 4 on a five-point scale.
    - Potential impact of the application received a score of 4.17 on a five-point scale.
    - Aesthetic appeal of the application received the highest score, 4.67 on a five-point scale. Several participants even mentioned during the course of the study that they especially liked the aesthetics of the visualizations presented.
    - Interest in using the application received a score of 4.17 on a five-point scale.
- The survey also gave them an option to leave feedback on overall impression of the application. A selection of responses is presented below:
    - "Pretty cool idea. Definitely brings improvement to the problem space."
    - "Positive!"
    - "Very useful, could definitely see it being deployed"
    - "Beautiful visual representations of online discussions"
    - "Useful for both in-depth reading and getting a quick overview of activity in a (discussion) space"

Several improvements that would enhance the functionality of the application were suggested by participants as well:

- More than one participant suggested that there should be functionality provided that highlights the current comment being read in the discussion overview visualization.
- Two-way moderation filter was another frequent request. While one slider filters out content, the other slider control would filter-in content. With this component, it would be possible to filter out only the mediocre comments,

while retaining both the well-received and controversial comments. This would also allow the user to filter down to the negatively-moderated/ controversial comments, something that is not possible with the current implementation.
- Another participant suggested that a visualization that highlighted the temporal placement of comments within a thread would be a beneficial addition to the suite of visualizations presented by application.
- Other improvement suggestions such as transitions for expansion of threads, transitions for opening the visual overview were mainly limitations of the implementation.

Another round of user testing, which will include participants who are actual users of Reddit, has been planned. All the usability problems uncovered in the preliminary round of user testing will be fixed before the next round, so as to enable feedback on the ideas developed without the detriment caused by minor usability issues.

Future Work

The area of large-scale discussion spaces is one that is teeming with interesting research questions. Discussions afforded by newer social media such as Reddit are notably different from other discussion spaces such as USENET, which have been the focus of researchers in the past. Further research into the dynamics of these conversations could be influential in driving the design of future interfaces for discussion spaces.

A study looking into the growth of a single discussion, and how variables such as growth rate, activity rate vary depending on the position of a thread in the discussion space could lead to several insights on user behavior.

By deploying the application on a larger scale, several interesting opportunities open up. Use of visualizations in interfaces meant for widespread use is something that has just started to gain traction on the web. A study looking into how these visualizations were received when deployed on a heavily used application such as discussion spaces would be beneficial. While this paper presents novel interfaces for navigating through discussions, considerations on how the interface could change the discussion is not discussed here. Deploying the application enables studies such as this to be conducted as well.

While each of the visualizations presented in the Data Analysis section of the paper highlight certain patterns within conversations, it would be more useful as an analytic tool with support for brushing and linking between the visualizations. An application built for analysis of data within discussions is currently being developed.

Acknowledgements

I would like to thank my advisor, Coye Cheshire, for guidance and encouragement throughout the project. I would also like to thank my fellow colleagues at the School of Information who have provided invaluable support throughout the project. Special thanks to the wonderful team at Reddit, who have provided much help with data access and the survey, and to the Reddit community for all the insight and interest.

REFERENCES

1. Sack W - *Discourse Diagrams: Interface Design for Very Large-Scale Conversations* - HCISS, 2000
2. Donath J - *A Semantic Approach to Visualizing Online Conversations* - Communications of the ACM. Apr 2002.
3. Smith M A, Fiore A T, *Visualization Components for Persistent Interactions* - CHI 2001.
4. Dave K, Wattenberg M, Muller M - *Flash Forums and ForumReader: Navigating a New Kind of. Large-Scale Online Discussion* - CSCW2004.
5. Dourish P, Chalmers M - *Running out of Space: Models of Information Navigation* - CHI 1994.
6. Whittaker S L et al. - *The dynamics of mass interaction*, ACM Conf. CSCW, 1998.
7. Bawden, D. 2001. *Information overload*. Library & Information Briefings 92. http://litc.sbu.ac.uk/publications/lframe.html
8. GK Zipf, Human behavior and the principle of least effort, Addison-Wesley Press 1949
9. Hiltz, S., and Turoff, M., *Structuring CMCS to avoid information overload*, Communications of the ACM 1985.
10. Jacob Jacoby, *Perspectives on Information Overload*, Journal of Consumer Research, Mar 1984.
11. Jones Q, Ravid G, Rafaeli S - *Information Overload and the Message Dynamics of Online Interaction Spaces: A Theoretical Model and Empirical Exploration* INFORMS 2004.
12. Schultz U, and Vandenbosch B, *Information Overload in a Groupware Environment: Now You See It, Now You Don't*, Journal of Organizational Computing and Electronic Commerce, 1998
13. Wattenberg M, and MIllen D, *Conversation thumbnails for large-scale discussions* - CHI 2003.
14. Flare Visualization Toolkit - http://flare.prefuse.org
15. Nielsen J, *Why You Only Need to Test with 5 Users* - http://www.useit.com/alertbox/20000319.html