Not too long to read: The tldr Interface for Exploring and Navigating Large-Scale Discussion Spaces

Srikanth Narayan
School of Information, UC Berkeley
srikanth@ischool.berkeley.edu

Coye Cheshire
School of Information, UC Berkeley
coye@ischool.berkeley.edu

Abstract

We present a new interface for exploring and navigating large-scale discussions on the internet. Our system, tldr, focuses on three primary user goals: finding, navigating and filtering content based on the users' own relevance criteria. We employ a progressive design approach that combines user research from the popular web site Reddit.com and visualizations from existing discussion data. Based on lessons gleaned from previous research and our own user studies, we argue that the tldr interface can lessen the problems of information overload and inefficient navigation in large-scale discussion spaces. Current implementation and next steps for wider deployment are discussed.

1. Introduction

The internet allows millions of people to easily exchange information and participate in numerous conversations with others around the world. Many different platforms currently exist on the internet to facilitate both long-term and fleeting conversations. Web forums, message boards, Usenet newsgroups, instant messaging, and email are just a few of the most popular systems for online discussions. While some systems such as instant messaging and email are primarily used for one-to-one exchanges, other online discussion spaces are used for shared public conversations between many different people at once. These large-scale conversations are asynchronous communications and thus enable interaction between participants without requiring simultaneous or continual presence among all participants [13].

Many large-scale discussions are user-initiated, user-maintained and user-moderated. A user can create a topic or present a question and reach a wide audience while keeping up with additional, ongoing conversations. Large-scale conversation systems foster discussions on every imaginable topic and bring in opinions from a variety of perspectives. In addition to the practical purpose of answering questions, these systems also encourage open-ended discussions between stable and ephemeral groups of users. Once viewed as revolutionary, the ubiquity of chat rooms and web boards over the past few decades has made these large-scale discussion spaces an ordinary and conventional way to communicate with others on the internet [19].

The number of large-scale discussion spaces continues to grow on the internet, in large part because many web destinations incorporate discussion components into their existing systems. Discussions on popular websites grow proportionally with traffic, and as a result, it is not uncommon to encounter discussions with hundreds or thousands of messages. This phenomenal growth can be observed on a wide variety of websites—news outlets, blogs, social media websites, community websites and support forums to name a few. While most of these discussion spaces are usually able to support smaller discussions, usability and manageability are greatly reduced as the discussions grow in size. Communication systems that support interactions in very large numbers (e.g., hundreds to thousands of unique participants) are quite different from those that support smaller conversations. Indeed, the dynamics of massive conversations are poorly understood due in large part to the complexity that comes with size.

Web-based discussion spaces have become a popular way to communicate on the internet yet they are technologies that have not seen consistent advances or improvements for user exploration and navigation. Currently, most websites present discussions as a linear list of messages that require users to scroll and read through tens of thousands of lines of text across multiple linked pages. Popular tools such as phpBB and vBulletin primarily manage large discussions by breaking content into increasingly smaller parts such as forums, topics, threads, posts, and pages. Few if any systems incorporate any type of visual interface to aid exploration and navigation.

IEEE computer society

978-0-7695-3869-3/10 $26.00

A consequence of compartmentalizing conversations into categories, hierarchies and threads is that it makes it difficult or impossible for a user to gain an overall view of where she is located within a given discussion space. Among the smallest units in the hierarchy are specific conversations—which may have many different branching trees that extend for thousands of lines of text or more. As individual discussions grow longer, it exacerbates the existing limitations of the interfaces. Users are left with a suboptimal dichotomy: they are either overwhelmed by the sheer amount of information presented at once, or they are offered micro-views of specific content with little or no connection to the larger conversation space. The chasm between viewing specific content and appreciating the surrounding discussion space would seem to implore designers to create informative visualizations of large-scale discussions for users. As Sack argues, an important part of successfully navigating large-scale discussions involves linking the "forest" of the discussion space to the "trees" of content [13].

In this paper we introduce tldr, an interface for exploring and navigating large-scale discussions on the internet. The interface is named for the internet slang term for extremely long threads, messages or stories, "tl;dr". Literally, "Too long; didn't read". In the context of a discussion space, tldr is used to indicate a summary for those who don't want to read the entire message.

We begin by contextualizing the problem of reading and traversing large-scale discussion spaces. Since different systems have their own design features and idiosyncrasies, we focus our efforts on a single large-scale discussion system, Reddit (Reddit.com). To better understand the landscape of expectations and behaviors on this system, we present the results from a large survey of current Reddit users. In addition, we use data from existing conversations on Reddit to define the scope of our visualization design process. Drawing from these two sources of data we show the primary user-focused design goals for the visualization interface and explain the key stages of our design process. Finally, we provide results from user testing and discuss the future directions for the tldr interface.

2. Large-Scale Discussion Spaces: The Dynamics of Asynchronous Online Conversations

The topic of large-scale discussions has been of interest to researchers since the emergence of the Usenet in the late 1970's. Usenet was the first internet system to organize topics and newsgroups into widely distributed hierarchies that support endless numbers of threaded conversations in publicly viewable conversations. The quick growth and ubiquity of Usenet was due, in large part, to the fact that it was a completely unregulated system for mass interaction [19]. Like email, Usenet was designed for asynchronous and persistent communication. A single conversation thread can last as little as a day, or as along as a year or more. By the 1990's the Usenet was, "the third most widely used form of interaction media on the Internet (behind email and the World Wide Web)" [15].

2.1. From Usenet to Web-based Flash Forums

Although the Usenet was the first large-scale public discussion space on the internet, its popularity as a platform for asynchronous conversations began to wane as web-based forums and message boards emerged. One of the biggest changes in the landscape of computer-mediated communication at the end of the 1990's was the increasing reliance on the web browser as a primary interface, ultimately providing a convenient and easy way to participate in online communication such as asynchronous discussion forums [6].

Most large-scale discussion spaces currently exist on the internet within specific websites as web forums, message boards and threaded comment systems. From the perspective of users, the specific technologies usually differ in terms of relatively marginal cosmetic features. Examples of some of the most popular large-scale discussion spaces include Reddit.com and Digg.com (user-maintained news aggregation websites), Slashdot.org (a site based on user-submitted stories with appointed moderators), NYTimes.com (discussion of stories published in the popular newspaper), and Engadget.com (blog of current gadget and technology news).

In their investigation of Slashdot forums, Dave, Wattenberg and Muller argue that many of the web-based forums that are becoming popular on the internet are significantly different from Usenet discussions [2]. The Usenet organizes disparate topics into categories and hierarchies into one system. However, web forums on sites with clear objectives such as Slashdot, Reddit, Digg, etc are located *within* these web sites. The result of this arrangement is that people who come to the main website can jump in and out of the forums whenever they like. Unlike most Usenet readers, this is done without leaving the browser interface. The primary characteristics of these quick and often fleeting web forum conversations are their diffuse authorship, large size, focused topics and constrained time frames [2].

2.2. Managing Information Overload

As discussion spaces grow, many problems emerge in proportion to the number of new content and users. These include an inability to easily find messages of interest, low signal-to-noise ratio, greater redundancy, and missed communication input. As the amount of information collected within a discussion space grows, the lack of ways to control and filter through the mass of information further aggravates these problems [1].

Research in cognitive science and psychology consistently shows that humans are governed by basic limits in our abilities to process information. For example, as discussions grow larger, the amount of information that merits attention can exceed the ability of the user's ability to process it [14]. Individuals who attempt to manage more information than they need or want can become overloaded, which in turn leads to inefficiencies for the information system and the user [11]. When confronted with too much information, users tend to resort to the path of least effort [21] and develop strategies to evade being overloaded [8]. For example, Jones, Ravid and Rafaeli found that information overload can lead to a marked decrease in participation. Worse, in some cases information overload can end active participation altogether [9].

In their work on information overload and the structure of computer-mediated interaction, Hiltz and Turoff suggest that, "The value of an information system lies in what it withholds, as much as in what it gives" [7]. If a system is not organized so that it can be interpreted in a reliable, coherent fashion, information overload is a likely consequence of this disorder. *Information entropy* refers to situations in which a lack of organization prevents users from detecting important content [7]. Of course, what constitutes 'relevant' content will differ between users. Thus, providing the user with tools for exploring and navigating large amounts of content while maintaining organizational coherency alleviates many of the problems created by large quantities of information. This is the guiding principle for tldr and the impetus for many of the related projects reviewed below.

3. Related Work: Visualizing Large-Scale Discussions

Several projects have addressed the issues of exploration and navigation in large-scale discussion spaces. Some of these projects are decidedly analytic in nature while others endeavor to create better end-user interfaces. Here, we briefly describe some of the key advances in visualizing large-scale discussions.

Given the aforementioned success of the Usenet over the past few decades, it is no surprise that it was the focal system for the first examinations of large-scale internet discussions. Sack's Conversation Map provides a visual representation of conversation threads in Usenet over time [13]. The visualization includes social network information and a semantic network of themes within the conversation. One of the key benefits of the conversation map is that it provides users with a tool for exploring the conversation, allowing one to bring up the content of individual messages from the visual interface. The visualizations are presented as the primary interface, and the content of the discussion are one layer beneath the visualizations.

The Loom project by Karahalios and Donath presents online discussions on Usenet through a semantically-based visualization [3]. A primary purpose of the project is to elucidate the limitless stream of statistics that can be culled from newsgroups into socially meaningful classifications. The Loom system explores how different visual representations can display meaningful and intuitive representations of the temporal distribution of a participant's contribution, patterns within sequence of messages, and patterns based on content of messages.

Smith and Fiore detail another system that highlights structural and temporal patterns within Usenet conversations using thread-based tree visualization of the social network within the conversation [16]. The primary goal of the project was to develop visualization components for threaded conversations that enable deeper understanding of patterns of activity. In addition, their thread-based tree visualization provides greatly enhanced visibility to the otherwise, "semantically meaningless hierarchy of tens of thousands of newsgroups" [16]. Other visualization systems for viewing threaded conversations include Venolia and Neustaedter's mixed-model visualization [17], and Kerr's Thread Arcs [10].

Yee and Hearst's content-centered discussion map [20] has a visual style that is helpful when viewing a specific thread within a larger set of threaded conversations. Color is used to highlight messages represented by the same user as 'blocks' in a tree diagram. In addition, content that is deemed especially important (such as unresolved questions in a thread) is emphasized through the user interface. Although the content-centered discussion map worked fairly well as a way of viewing previous conversations, the system did not transition to a real-time visualization and navigation interface for ongoing discussion spaces [20]. Newman's TreeTable explores a similar visualization technique, developed to obtain thread

overviews and mechanisms to aid in coherent reading of threads [12].

Dave, Wattenberg and Muller focus on large-scale, topic-centered, transient discussions called 'flash forums' [2]. Using data from the popular news aggregator Slashdot, they present the ForumReader 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 long-term persistent threaded conversations (e.g., Usenet) compared to more transient discussions such as Slashdot. For example, in their experimental test of features in the ForumReader system the researchers found that author information was much less important than previous research on long-term threaded conversations would suggest [2].

Discussions afforded by newer social media websites such as Reddit are considerably different from the conversations on Usenet, and are arguably much more closely related to the flash forums studied by Dave, Wattenberg and Muller [2]. Though tldr and ForumReader have similar goals, our motivation was not to simply improve upon systems like ForumReader. Instead, we focus on user needs in a specific online system, and design a visual and navigational interface to match these needs.

In summary, much of the previous research on large-scale discussions has dealt with Usenet and mailing list conversations. More recent work has begun to examine threaded web-forums and discussion groups. An important point is that many prior research efforts primarily concentrated on analytic interfaces and experimental visualizations, rather than interfaces specifically designed for end-users. The tldr interface presented in this paper differs significantly from many of the earlier attempts in this respect: the goal of this project is to propose and test designs that can be integrated into an existing system for users of a specific large-scale discussion space.

4. Exploring Reddit as a Large-Scale Discussion Space

Inefficiencies in handling large-scale discussions for users can erode the potential usefulness and value of a given discussion space. However, no single solution could address all problems across every discussion space. To better scope our efforts we chose a popular online discussion space, Reddit.com, as a model to investigate the dynamics of large-scale conversations.

Reddit is a popular social news aggregator web application with about 5 million unique visitors per month, and between 125,000-150,000 registered users. (Data obtained from http://www.compete.com, June 12, 2009). Users of this website submit web articles that might be of interest to the audience of the site. Articles that are popular among its users are promoted to the main front page. Each submission has an associated 'comments' component where users hold discussions about the topic of the given story. As a result of high activity on the front page, it is not uncommon for each of these discussions to reach hundreds to thousands of comments.

The current Reddit interface is extremely simple and primarily text-based. Despite being a usable, streamlined interface, it suffers from many of the limitations in handling large-scale discussions. Comment threads are presented in a linear expanded list and navigating through the interface requires an inordinate amount of scrolling. For example, a single discussion with 586 messages will be presented on a webpage that is 64,294 pixels long, or about 81 full screens long on a 1280x800 display. In addition to the inefficiencies of scrolling through text for the user, displaying the entire discussion at once also results in slower page loads.

4.1. Survey of Reddit Users

To gain a better understanding of users' perceptions and behavior on discussion spaces, we conducted a survey of Reddit users. Through a series of self-report questions, we collected information about individual levels of online participation, activities, interests and opinions about online discussions spaces. In cooperation with the administrators at Reddit, we ran a front page announcement linking to the survey. Over a fifteen-day period, we collected 477 valid responses. We offered no monetary incentive for participation. Thus, our sample is comprised of individuals who visited Reddit during a 2-week period and who were interested in contributing to research on discussion spaces.

We asked users to estimate the extent to which they participate in various activities within large-scale discussion spaces on a 5-point scale (1=never, 2=rarely, 3=occasionally, 4=often, and 5=very often). To get an overview of what people do most often, we examined those who said that they participated in an activity often or very often. A clear majority of the respondents (91%) indicate that they read comments often or very often. We found that 37% of users respond to comment threads and 31% moderate comments, but only 13% self-report that they start threads often or very often.

We asked survey respondents to assess the cues that aid them in navigating through large-scale discussion spaces. Table 1 indicates that the user's

own perception of the message after skimming the first few sentences is a primary way that one navigates through a discussion. Position in the discussion space was also found to be very important, implying that thread-relevance algorithms are implicitly accepted as useful by many users. Moderation score was also found to be highly important, indicating that users place a high degree of trust in their peer community to find and rate significant messages within the discussion. The author of a given comment was found to be the *least* effective cue. In large online discussions, it is arguably very hard for an individual poster to stand out. The relative unimportance of the author complements Dave, Wattenberg and Muller's similar finding for the flash forums on Slashdot.org [2].

	Very Unimportant	Somewhat Unimportant	Neither imp. nor unimp.	Somewhat Important	Very Important
Moderation Score	7.3%	14.6%	13.4%	47.7%	16.9%
Comment Author	48.8%	15.9%	18.4%	13.4%	3.4%
Comment Length	4.6%	16.6%	28.6%	43.5%	6.7%
Number of Replies	7.4%	20.6%	24.2%	37.6%	10.3%
Position in Discussion	4.2%	7.1%	13.4%	38.7%	36.6%
Message Perception	2.3%	4.4%	8.6%	43.9%	40.8%

Table 1. Percent responses for cues that aid navigation in discussion spaces.

We also asked a single yes/no question about whether the respondent ever uses filters to limit the number of comments in a discussion space. The majority of respondents (60%) indicated that they *do not* use any type of filter. We asked those who self-report that they do apply filters to list which one(s) they use: collapsing comment threads, limiting by moderation score, limiting by number of messages, or other method. Among respondents who said that they use filters, 52% listed 'moderation score limit', 55% 'collapse by thread' and 39% 'number of messages'.

4.2. Exploring Data from Reddit Discussions

In addition to our user survey research, we also analyzed existing data within large-scale discussions on Reddit. These efforts were primarily focused on understanding data that comes from within individual large-scale discussions and applying visualizations that could aid in exploration and discovery.

In the first visualization shown below (Figure 1), each vertical column represents an hour. All of the comments that fall into a given hour are stacked in the same column. 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 large numbers in a short duration. Color saturation by positive/negative moderation 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. A gradual fading of color over time hints that it is often hard for a comment worthy of merit to be recognized as important unless it is posted earlier.

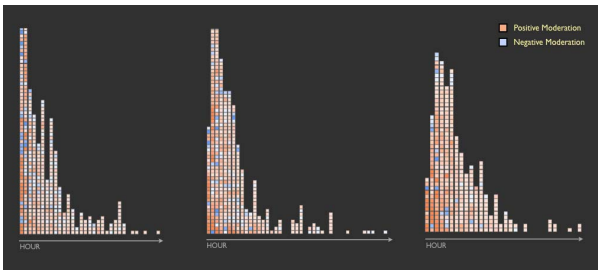

Figure 1. Histogram visualization of comments (3 random discussions).

The next visualization (Figure 2) was aimed at uncovering patterns within conversational exchanges in a single discussion. While it would be an arduous task to estimate the quality of all messages by analyzing their content, patterns in moderation score can be rough indicators of quality 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. Thread 2 seems to have a broad ongoing discussion. Certain branches of the tree tend to be favored higher by readers than other branches (a), and certain branches of the tree grow deeper than other branches (b).

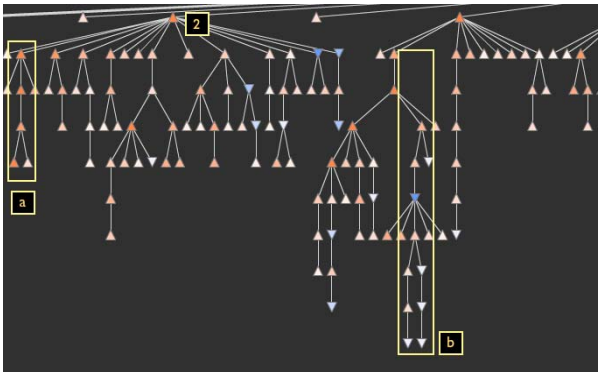

Figure 2. Tree diagram of a discussion (Noteworthy patterns highlighted).

The patterns shown above could be made more conspicuous by summarizing individual branches at the same conversation depth. The next visualization (Figure 3) is based on an icicle tree layout in which

nodes of unit length are 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.

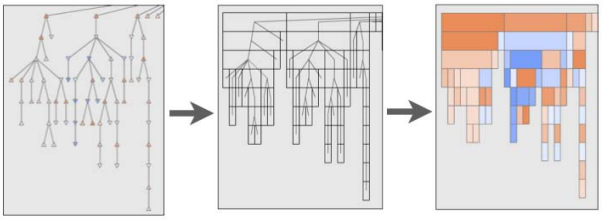

Figure 3. Transforming the tree to make patterns conspicuous.

Finally, the last visualization (Figure 4) was designed to facilitate comparison across threads. Activity, moderation trends, length of comments, and level of reply are important variables that aid in comparing threads. The visualization in Figure 4 builds upon the original design by Wattenberg [18], while enhancing different patterns within a discussion. Specifically, length and activity are easy to compare across threads. The length and activity of individual comments within a single thread are also easy to evaluate. For example, the first two threads in Figure 4 (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). The progression of activity across threads can also be seen, hinting at how the placement of threads within the discussion space affects relative activity rates.

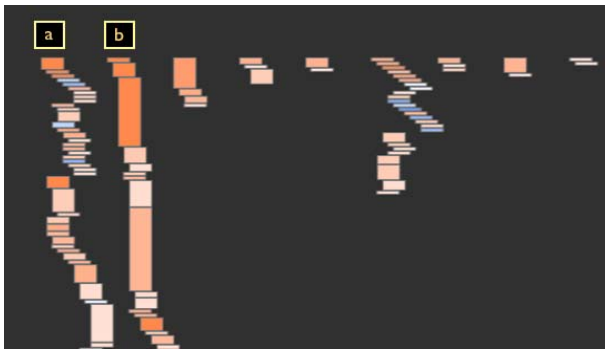

Figure 4. Modified thread comparison visualization (based on [18]).

5. tldr: A Novel Interface for Large-Scale Discussions

In this section we present the tldr interface and visualization application. Based on our user survey and visualization prototypes, we identified three key user goals in large-scale discussions: identification, navigation, and filtering of interesting content.

Identification of content of interest. Few users would want to read through entire discussions with thousands of messages. While many discussion spaces incorporate functionalities for users to assess messages of interest, the specific cues that help individuals make such judgments are not always presented effectively. Thus, a visual interface should aim to let users make qualitative assessments of interesting content as easily and as transparently as possible.

Navigation to content of interest. Given the large number of messages in a discussion, content of interest to a particular user may be scattered across the discussion. Common linear or paginated lists are ineffective for navigating in this situation. Keyword searching is the most common solution, but this method tends to trade discussion context and position in favor of content matching. An interface and visualization system should allow individuals to find content without losing awareness of position within the discussion space.

Filtering to content of interest. In a discussion with thousands of messages, users may wish to limit the number of messages displayed. However, the process of filtering out messages can inadvertently reduce one's awareness of the surrounding conversation space. A visualization and interface should aim to let the user filter messages along several dimensions while maintaining an accurate representation of the surrounding discussion space.

5.1. User Interaction with tldr: Addressing the Three Key User Goals

Figure 5 illustrates the primary interface that is presented to the user. In the primary view, the discussion is presented at the thread level and all ongoing threads are laid out so that the user can get a sense of activity within the surrounding threads. This view provides relevant information about each thread, enabling the user to quickly browse across threads and dive into different parts of the discussion.

The primary interface also provides an option to open a visual overview in the upper right-hand corner (see Figure 6). This 'Discussion Overview' presents a visualization of the entire discussion, highlighting the

patterns within these ongoing conversations. As a result, the visual overview facilitates exploration at the discussion level (one level higher than the thread level presented in the previous view).

Figure 5. The primary tldr interface.

The visualizations presented in the interface not only serve to highlight patterns within the discussion, but also act as a navigational aid for the user. Interactions such as hover over the nodes of the visualization lets the user learn more about the posted messages, and click on any node within the visualization lets the user navigate directly to that message. As a result, the interface presents a non-linear navigation scheme that enables skipping directly to messages of interest across an entire discussion space.

In addition to the functions described above, the application also allows filtering of messages in the discussion along several dimensions (moderation level, semantic tags and activity level). Interacting with these filters updates the visualizations and the textual display, thus controlling the number of messages presented in the interface.

5.1.1. Identification of Interesting Content

Thread Summarization. The initial view of the interface presents summaries of activity within threads. Each one of the blocks is an ongoing thread, and relevant information about the thread is presented within the display. The visualization is presented in the context of each thread and allows the user to infer popularity, length and overall structure of the

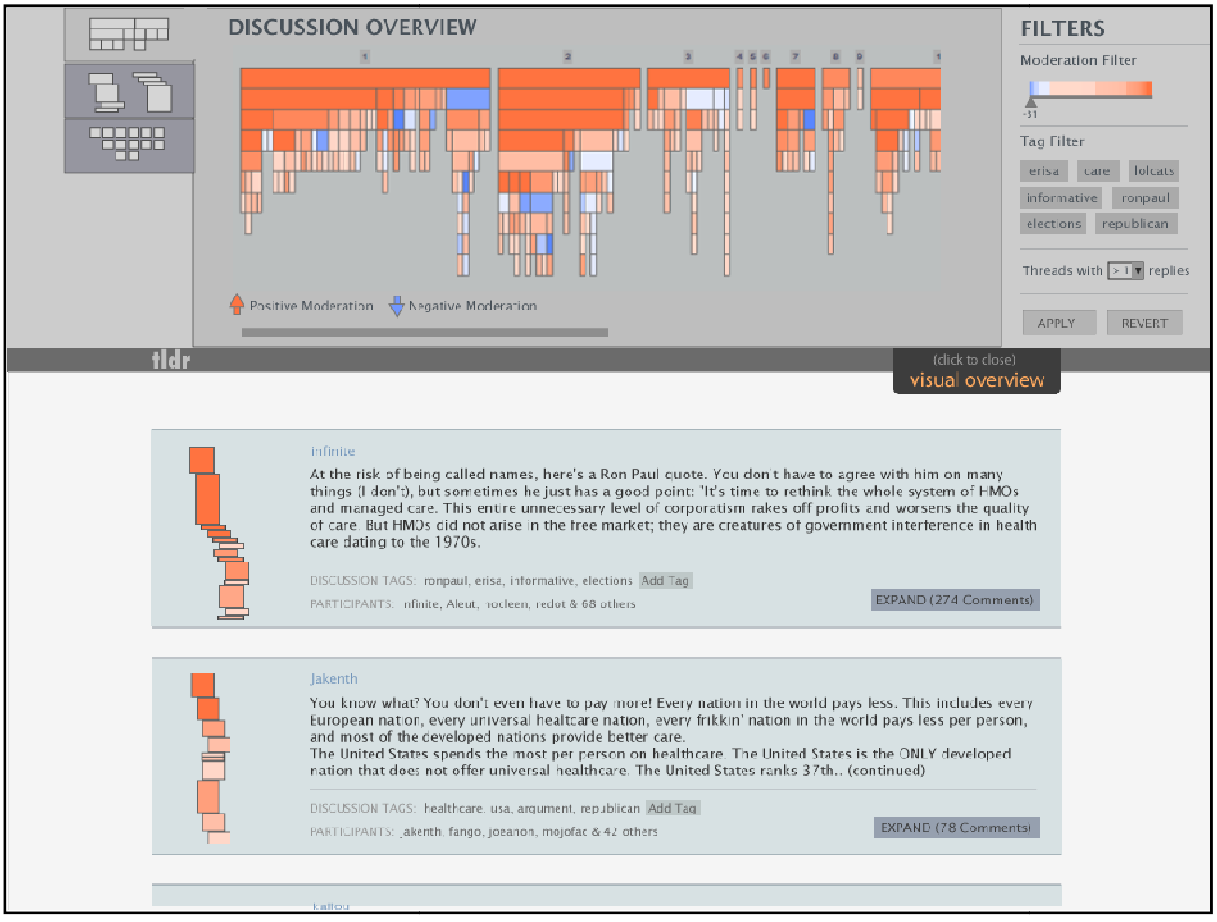

Figure 6. tldr interface with *Discussion Overview* open.

conversation. This supplements the information that is provided in the thread overview and helps the user make a more qualified assessment of the worthiness of a conversation thread. Upon finding a thread of interest, the user can dive into messages within that thread by 'expanding' the overview. The overview is then replaced with an expanded list of messages presented in an indented-tree fashion (see Figure 7).

Figure 7. The Thread Summarization view.

Discussion Summarization. The 'discussion overview' view presents several visualizations that depict activity within an entire discussion (see Figure 8). These visualizations not only make patterns within conversations apparent to the user, but also serve as a navigational aid to travel into other points of interest. This view also presents several filters that enable users to limit the messages along different dimensions. The discussion overview is considered a supplement to the core functionality provided by the application. Thus, the overview panel is hidden by default. When the user wants an overall picture of the discussion to navigate to points of interest, she can invoke this view.

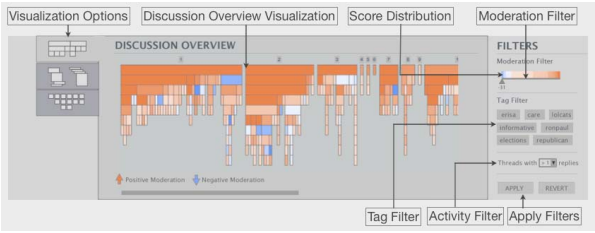

Figure 8. Features of the discussion overview.

In the discussion overview visualization, each block represents an ongoing thread in the discussion. The arrangement of nodes is such that it conveys information about the position, reply level and popularity of the message that it represents. Studying the visualization makes it easy to discern threads with high participation, that run the deepest, conversations that have been favored the most, controversial areas of the conversation, overall distribution of activity across threads, and individual comments within threads that are anomalistic (For e.g., a witty comment that has gained many more votes than its surrounding comments). The user can also switch between visualizations in this view, rendering the same data to gain different insights into the discussion.

5.1.2. Navigation to Content of Interest. Upon finding something of interest within the visualization such as an active area or highly rated series of comments, the user can click on the nodes within the visualization to navigate directly to the points of interest. Clicking on a first-level node within a thread leads the user to the overview for that thread. Clicking on any particular node within the visualization leads the user to the message that the node represents. The message is presented in the context of the conversation that it was part of, so that it is easy for the user to read through the entire conversation. Other messages within the thread that were not part of this conversation are collapsed, thus enabling reading through a coherent conversation without much conversation drift [5]. This navigation scheme also encourages users to explore the discussion. The primary purpose of the navigation aspects of the interface is to make it easier and faster to move across threads. In addition, the navigation tools increase the exposure of content and allow for serendipitous discovery of interesting content.

5.1.3. Filtering to Content of Interest. Filtering is an important part of finding relevant content, but our user survey indicated that this is often an under-utilized feature in large-scale discussions, perhaps due to the inaccessibility of these features on Reddit. Our design goal for addressing this issue was to limit the display of content while maintaining an accurate visual representation of the thread and/or discussion. Thus, we employ filters that change the user view based on the selected criteria, as well as 'highlight' filters that emphasize content that match filter information such as semantic tags. The filters seen in Figure 8 allow the user to limit the number of messages displayed in the visualization and the textual display along three different dimensions:

Moderation filter: allows the user to limit the comments displayed based on the moderation score it has received.

Semantic Tag Filter: allows the user to select comments within a discussion that have been tagged with a particular qualitative descriptor. Hovering the cursor over any particular tag highlights all the messages that have been tagged with that descriptor in the visualization. A user can make multiple tag selections to highlight messages of interest. Further, the user can apply these selections to the textual display to expand all the conversations that are tagged with the descriptors chosen.

Activity Filter: allows the user to limit threads displayed based on the number of messages in the thread, enabling users to easily ignore mediocre messages that have not received much attention, or directly filter down to threads that have received the most attention.

6. User Evaluation of tldr

After developing a working prototype, we conducted a small user study to evaluate the efficacy of the tldr interface. First, the participants were asked to explore the functionality offered by the application. A think-aloud protocol was employed to gain an understanding of how participants perceived the functionality presented by the interface. Next, participants were asked to perform several scenario-based tasks that tested their understanding of specific features (e.g, the user's ability to identify certain messages in a large thread, identify threads with lengthier postings compared to others, etc). Finally, participants filled out a short exit survey.

Six participants agreed to participate in the user test which was conducted at a large research university. All the participants were regular to heavy users of large discussion spaces including Slashdot, deviantART, Digg, etc. Each study took approximately 45 minutes.

The application was well-received by our users. Most participants were fairly interested in using the application, and many explored the tool beyond the requirements of the task. Among the visualizations presented, participants mentioned that they found the icicle visualization compelling because it enabled recognition of potentially interesting messages and conversations, while facilitating quicker navigation through the entire discussion. Several participants positively commented on the design decision to hide the Discussion Overview by default, explaining that it incrementally provided more features as one explored the interface, without overwhelming the user with options from the beginning. The majority of participants were able to understand all of the visualizations and functionality presented in the application without any guidance from the moderator.

The exit survey asked participants to rate the effectiveness of various features on a five-point scale from 'Not at All Effective' to 'Extremely Effective' Features evaluated were: Thread Overview, Thread Thumbnail Visualization, Discussion Overview Visualization, Navigation using Discussion Visualization, Moderation Filter and Tag Filter. Participants also rated different how well the interface accomplished certain design goals: Usefulness of the application, Ease-of-use, Potential impact of the application, Aesthetic appeal of the application and Interest in using the application. Quantitative results from such a small user study should be interpreted with care as this is not a large, representative sample. Still, it is encouraging that all of the features were rated effective or highly effective by the participants.

The open-ended user feedback included positive comments concerning the look and feel of the visual representation, usefulness of the interface for in-depth reading, and the ability to observe a quick overview of activity. Overall, the feedback indicated strong interest in using the tldr interface for navigating through large discussions. More importantly, the user study was crucial for uncovering usability problems with the application and guiding future iterations of the design. Transitioning from the Thread Summary view to the textual representation of the thread was seen as confusing, as the views replaced each other. Another concern was the heavy reliance on scrolling, which is not a comfortable mode of interaction for some users.

Participants were proactive in making suggestions for further improvement. The most requested feature was the ability of the visualization to highlight the specific comment or thread currently being read. In conjunction with the existing navigation features, this would make it easier for users to visually track their position as they move through a discussion space. Another frequent request was a two-way moderation filter that would allow a user to specify a range of moderation scores. Finally, some participants suggested visualizations that would highlight the temporal placement of messages in a given thread.

7. Discussion and Future Work

Our design strategy for the tldr interface demonstrates the importance of user-focused research for creating effective visualizations in an online system. A single visualization that can apply to all systems and users is impractical. However, we submit that the tldr interface is a useful and effective tool for users of the text-based threaded discussions in the Reddit forums. Our focus on the needs of real users aided our development of a practical and intuitive system.

The next steps in the tldr design process include the deployment of the application on a larger scale and increased user customization. Visualizations in discussion interfaces are relatively nascent on the web. While this paper presents novel interfaces for navigating through an existing discussion space, investigating how the interface could *change* the exploration and navigation activities of current users is another important goal.

The area of large-scale discussion spaces is one that is teeming with interesting research questions and

opportunities for interfaces that improve exploration and navigation. For example, studies that examine the progression of a single discussion and how variables such as growth and activity rate vary depending on the position of a thread in the discussion space could lead to several insights on user behavior.

Online interaction spaces are constantly evolving. Designers of visualizations and interfaces must consider the nature of the information that is presented in these spaces, as well as the spatial, semantic and social navigation features that they engender (cf. [4]). Further research into the dynamics of visual presentation, search and navigation in large-scale discussion spaces is essential.

8. Acknowledgements

We thank Andrew Fiore, Marc Smith, Marti Hearst, Judd Antin and our p for their invaluable feedback. We especially thank the founders of Reddit for their cooperation and support in our survey.

9. References

[1] Bawden, D. Information Overload. *Library & Information Briefings 92* pp. 1-15 (2001)

[2] Dave, Kushal, Martin Wattenberg, and Michael Muller. "Flash forums and forumReader: navigating a new kind of large-scale online discussion." *CSCW '04: Proceedings of the 2004 ACM conference on Computer supported cooperative work.* New York, NY, USA: ACM Press, 2004, 232-241.

[3] Donath, Judith. "A semantic approach to visualizing online conversations." *Communications of the ACM 45* (April 2002): 45-49.

[4] Dourish, P. and M. Chalmers. "Running Out of Space: Models of information navigation." *HCI, Glasgow* (1994).

[5] Herring, Susan C. "Interactional Coherence in CMC." the *32nd Annual Hawaii International Conference on System Sciences* (1999).

[6] Herring, Susan C. "Slouching Toward the Ordinary: Current Trends in Computer-Mediated Communication." *New Media Society 6* (February 2004): 26-36

[7] Hiltz, Starr R. and Murray Turoff. "Structuring computer-mediated communication systems to avoid information overload." *Communications of the ACM 28* (July 1985): 680-689.

[8] Jacoby, Jacob. Perspectives on Information Overload. *Journal of Consumer Research*. Mar 1984.

[9] Jones, Quentin, Gilad Ravid, and Sheizaf Rafaeli. "Information Overload and the Message Dynamics of Online Interaction Spaces: A Theoretical Model and Empirical Exploration." *Info. Sys. Research 15* (June 2004): 194-210.

[10] Kerr, Bernard. Thread Arcs: An Email Thread Visualization. *Proceedings of IEEE Information Visualization* (2003): 211-218.

[11] Losee, R. Minimizing information overload: The ranking of electronic messages. *Journal of Information Science. 15,* 3 (1989), 179–189.

[12] Newman, P. S. Exploring discussion lists: steps and directions. *Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries.* (2002): 126—134.

[13] Sack, W. "Discourse diagrams: interface design for very large-scale conversations." the *33rd Annual Hawaii International Conference on System Sciences*. 2000, 9-10.

[14] Schultz, U. and B. Vandenbosch, Information Overload in a Groupware Environment: Now You See It, Now You Don't. *Journal of Organizational Computing and Electronic Commerce*, Vol. 8, No. 2. (1998), pp. 127-148.

[15] Smith, M. A. Invisible crowds in cyberspace: measuring and mapping the social structure of USENET. In M. A. Smith & P. Kollock (Eds.), *Communities in Cyberspace: Perspective on New Forms of Social Organization* (1999):pp 195-219. London: Routledge Press.

[16] Smith, Marc A. and Andrew T. Fiore. "Visualization components for persistent conversations." *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems.* New York, NY, USA: ACM Press, 2001, 136-143.

[17] Venolia, Gina D. and Carman Neustaedter. "Understanding sequence and reply relationships within email conversations: a mixed-model visualization." *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems.* New York, NY, USA: ACM Press, 2003, 361-368.

[18] Wattenberg, Martin and David Millen. "Conversation thumbnails for large-scale discussions." *CHI '03: CHI '03 extended abstracts on Human factors in computing systems.* New York, NY, USA: ACM Press, 2003, 742-743.

[19] Whittaker, Steve, Loren Terveen, Will Hill, and Lynn Cherny. "The dynamics of mass interaction." *CSCW '98: Proceedings of the 1998 ACM conference on Computer supported cooperative work.* New York, NY, USA: ACM Press, 1998, 257-264.

[20] Yee, K-P, and Marti Hearst. Content-Centered Discussion Mapping, *Online Deliberation 2005 / DIAC-2005*, Palo Alto, CA, June 2005.

[21] Zipf, George. *Human Behavior and the Principle of Least Effort.* 1949.