Developing Legible Visualizations for Online Social Spaces

danah boyd, Hyun-Yeul Lee, Daniel Ramage, Judith Donath
{danah, hyun, dramage, judith}@media.mit.edu
Sociable Media Group, MIT Media Lab
20 Ames Street, E15-443; Cambridge MA 02139
http://smg.media.mit.edu/projects/loom2/

Abstract

Although constructed for researchers to share news and information, Usenet quickly developed into a social environment with varied styles of interactions. Unfortunately, the browsers developed to view the shared messages fail to effectively convey the rich social features of a newsgroup, let alone all of Usenet. The goal of our research is to use the salient features of social interaction to build a "legible" interactive visual representation of Usenet. In this paper, we introduce our approach to developing this type of visualization, discussing our theoretical framework, questions considered to access the socially salient features, and a series of design iterations used for exploring how to develop a visual language that conveys social meaning. Although this paper represents a work in progress, we hope that this approach and our initial iterations help build a framework for future directions.

1. Introduction

The city, however, does not tell its past, but contains it like the lines of a hand, written in the corners of the streets, the gratings of the windows, the banisters of the steps, the antennae of the lightening rods, the poles of the flags, every segment marked in turn with scratches, indentations, scrolls. [2]

Walking down a street in most cities, you are constantly given various physical and social cues that help you know how to interact with the space and the people around you. By briefly observing the people, their interactions, and the physical features of the space, you can quickly determine whether or not the environment is of interest to you.

Unfortunately, the same does not apply online. While online communities produce a vast amount of data about social interaction, this data is often impenetrable to the users. Nothing about the format differentiates a support newsgroup about HIV from a discussion group about the Simpsons TV show. Disentangling the style of interaction and developing an understanding of the people involved requires extensive reading and analysis, and does not necessarily help people recognize important social cues.

The online social environment of Usenet does not have a physical structure; the people and their interactions construct the space. In order to develop a sense of Usenet, it is crucial to understand the nature of the space, which includes the history of its people, their behaviors, and the content and intention of their messages. We are interested in developing a way to make these constructions legible for users so that it is easier to understand the social space upon initial inspection. By legible, we refer to Lynch's idea that the environment should provide contextual cues as to what the space is about, including its social norms [12].

In this paper, we start by introducing Usenet and other visualization research in this area, emphasizing our goals in developing a visualization of Usenet. Following this, we introduce which socially salient features of Usenet interest us and how they can be analyzed. Our next section introduces the design philosophies we consider when developing our visualizations. The remainder of the paper is devoted to discussing how we use these design philosophies to visually convey the aforementioned social attributes, focusing on portraying individuals, conversations and groups.

2. Why visualize Usenet, and with what approach?

As one of the oldest network applications, Usenet has evolved in the last 20+ years to provide a forum for all types of exchanges, ranging from technical question & answer groups to support groups, MP3 and pornography distribution hubs to political debates. Embedded in immense quantities of data are flourishing communities engaging in a range of social interaction. Unfortunately, the social characteristics are easily obscured by typical Usenet browsers, which emphasize the informative aspect of Usenet (Figure 1). In order to navigate the

conversations in these browsers, you must first select a group by finding its name and then choose a thread based on its subject. This approach lets you find content effectively, but isn't helpful in conveying the social qualities of the interaction. Given the innumerable styles of interaction in Usenet, immense possibilities exist exploring alternative techniques and approaches to visually convey social interaction.

Figure 1. Typical Usenet browser (Microsoft's Outlook Express)

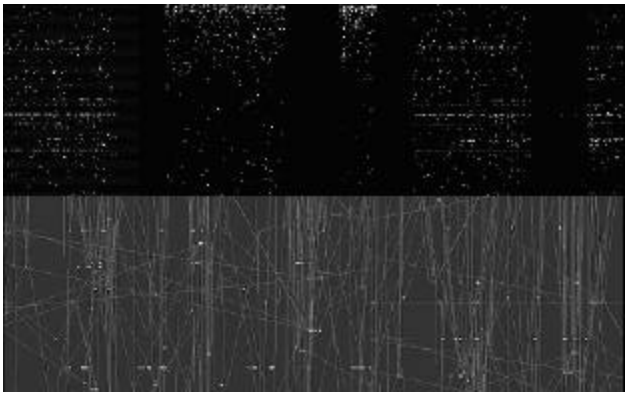

Figure 2. Loom – data visualization of Usenet, by Karrie Karahalios [6]

Many researchers, recognizing the potential of Usenet, have developed different mechanisms for visualizing important features of the environment. The original Loom [6], from which our research evolved, focused on using visualization techniques to uncover social patterns, ranging from pinpointing vociferous users to illuminating thread paths (Figure 2). Coming from a linguistics perspective, Sack [18] uses threading structure and language to construct "discourse diagrams," which convey the social and semantic networks present in Usenet. More recently, Smith & Fiore [21] have focused on portraying the social aspects of Usenet by visually expressing the quantitative qualities of the environment, such as number of people and frequency of posting (Figure 3). Alongside other research in visualizing online social environments and patterns (e.g., [7][8]), these three projects provide a foundation for visualizing Usenet's social data.

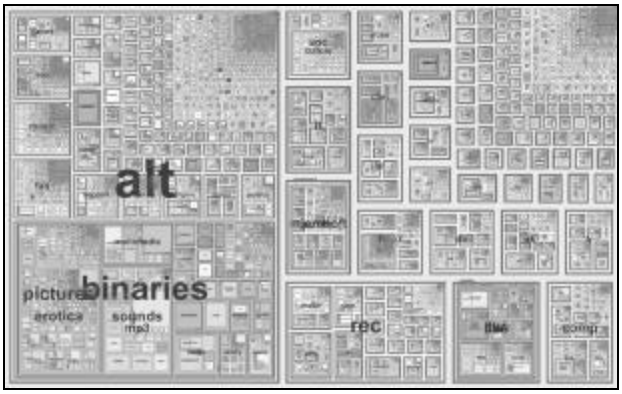

Figure 3. Netscan's box plot, landscape view of Usenet [21]

While detecting tone and emotional aspects out of text is still quite challenging, large quantities of social data are currently accessible. One current challenge is to determine how this data relates to the more fundamental social qualities of the environment. Rather than just considering the meaning of one piece of data, we must think about how they work together in a system. Another challenge is presenting not only the quantitative data, but also the underlying meanings, or social qualities.

Our visualization approach comes from both of these directions. Given a set of interesting social characteristics, how do we systematically analyze the data to reveal these features? Additionally, what visual attributes can be used to convey the social attributes of Usenet? While we have yet to build a bridge between these two approaches, we believe that our approach offers numerous opportunities for building a coherent visualization.

In some ways, this approach diverges from most information visualization work. Often, researchers working in information visualization are interested in developing techniques that will work for a broad range of data sets (e.g., [3][4][17]). Although we value the strengths of this approach, particularly in making quantitative information understandable, we also feel that it is limited when trying to convey the underlying qualitative nature of a particular environment. To represent a specific environment in this way, we must move beyond initial visualization techniques, such as larger visuals represents more in data [23], and develop ways to represent the more expressive and subjective characteristics.

In order to evoke a user's qualitative understanding, we need to use images that convey both the quantifiable and

the subjective aspects of the data. Detailed representations of each numerical attribute is less important than reflecting the fundamental differences between groups.

Relating subjective visual characteristics to interesting data must be intimately tied to the data that is being represented; in our case, we focus on large-scale textual social communication. More specifically, we are tied to the specific qualities of Usenet (i.e. how its structure and asynchronous nature affects social interaction). For this reason, we aren't focused on building visualization techniques that might work universally, although much of what we build does apply to other social visualizations, particularly those that emphasize group social interaction. Instead, our focus is in determining the socially salient characteristics and relating visual components to this information, bringing the underlying attributes, social patterns, and attitudes to the forefront of the visualization.

3. Recognizing social characteristics in textual data

Although the interaction in Usenet is entirely textual, many social characteristics emerge from it. For example, through regular use, one probably has a sense of the number of regular posters in a group and what type of tone they are trying to convey in their messages. Although it is possible to develop a sense of this through regular interaction, it is not easy to get the type of brief social overview upon entering a group as you would if you entered a physical environment, like a dance hall. What social qualities are important to one's initial impressions of a dance hall, and how is the same information conveyed textually in a digital environment like a newsgroup?

Focusing on the quantitative aspect of Usenet interaction, Whittaker [25] provides a series of quantifiable metrics that he believes contribute to mass interaction in a newsgroup; these are broken down into three categories – demographics, conversant strategies, and interactivity. Smith & Fiore [21] use many of the same metrics in their data collection, as well as more individual-centric features, such as interpersonal connections.

Including these quantitative characteristics as well as more qualitative ones, we developed the following series of questions that we feel convey the social characteristics of individuals, conversations, and newsgroups.

Individual. How frequently does an individual post? How verbose is s/he? What is the reputation of an individual? How many and which groups is s/he involved with? What can be said about an individual's identity, culture, job, or involvement in subgroups? What types of tone, language and conversation techniques are often used?

Conversation/thread. How many participants are active in a conversation? What structure or pattern of communication occurs (i.e. back and forth vs. large participation vs. one person dominating)? What is the tone of the conversation? What is the typical response to a newcomer's posts? Are messages posted here also cross-posted in other groups? How do the conversation patterns change over time or through extensive replies? Does the thread fracture into mini-threads? Are there common language patterns and terms?

Group. How many threads are usually active? Are there tonal differences between ongoing conversations? How many people participate in the group, and how frequently do they post? How many people start conversations? Is group activity consistent or do internal/external events trigger changes in posting patterns? What is the purpose of the group (i.e. support, Q&A, data sharing)? Do the participants have attributes in common, such as linguistic patterns, interests or group memberships? Are people who initiate threads also tend to reply to others' messages? What networks of people usually interact with one another?

Given these questions, it is important to imagine how to quantify them. Many of these questions are statistically easy to determine (i.e. How many? How often?). Others can be answered through known methods such as social network analysis [24] and various linguistics techniques, including semantic analysis [15]. The most challenging questions concern tone, not only to determine emotional content, but also to answer questions such as the main purpose of a group. Related work on flaming [22] provides one mechanism to analyze tone, but it is rather limited. Unfortunately, given the large quantity of Usenet data, even determining the statistically simple questions is computationally time intensive. Also, computing this information is memory-intensive process and requires local storage of large quantities of data. While these hard problems present significant challenges, for the purpose of this paper, we will assume that quantifiable data is readily accessible.

Given that we can quantify specific social attributes, what do these numbers actually mean? For example, if we know that a group has a large number of people interacting regularly, what can we say about the group? Is it vibrant or overcrowded? If a poster constructs long messages, is s/he conveying a detailed answer or ranting?

Some of these elements are personally subjective. For example, what may appear vibrant to us might seem overcrowded to you. In other cases, the answer is somewhat universally perceived; users reading a long message would probably recognize if a message were helpful to the conversation, or adding flame to the fire. Yet, in order to construct a dynamic visualization, we must systematically uncover such qualitative aspects. One

way to do this is to rely on the context of other information. For example, if the responses contain "angry" words, the post is most likely a rant or otherwise insulting.

Although understanding the meaning behind different characteristics helps direct the designs, we do not want to taint the view with only our personal perspective. Instead, we want to portray the information in a subjective, yet open manner, so that the users can determine the characteristics themselves. For example, we do not want to imply that a group with many conversations is vibrant or overcrowded; we only want to convey that there are many ongoing conversations. We try to minimize our personal biases while still conveying qualitative characteristics of the space, a delicate task.

4. Design fundamentals

Building a visual language with which to convey social information is a subjective process. Not only do we make assumptions about what social characteristics represent, but we also rely on assumptions about the effect that a visual form will have on the viewer. It becomes tremendously important to constantly evaluate the consequences of these assumptions.

In order to understand the impact of visual elements, we attempt to build on concepts developed by such fields as urban planning, graphic design, and interaction design. In this section, we will introduce these underlying philosophies. Three themes build the backbone for our research: Lynch's architectural theories [12] on legibility and related work on digital information landscapes; design and conceptual work on multi-scaled information visualization; and research done in the area of kinetic typography.

These building blocks create a conceptual framework for imagining how a user can interact with social information in a more accessible manner. By applying visual building blocks, such as shape, form, typography, color, and motion, to these philosophies, we hope to evolve a visual language that conveys the expressions of a social environment like Usenet.

4.1. Creating legible social landscapes

Legibility, or "the ease with which people can understand the layout of a place," is a fundamental of urban design; Lynch's work sought to create legible cities where shape, color, and arrangement facilitate people to create vividly identified, powerfully structured and highly useful mental images of the environment [1][12]. Although intended for the physical world, these concepts are easily translated to the presentation of "information landscapes," a term developed by Muriel Cooper [5] to describe a digital environment where information is geographically positioned in a virtual space. Just as urban designers must consider how to make a physical space legible, an information visualization designer must make the social landscape legible, thus allowing users to understand the space, its history, and its people through the visual representation of the environment.

In order for an environment to be legible, the design must rely on vividly identifiable components, thus implying that viewers share a certain level of cultural and cognitive intuition about the presentation. Although Raskin [14] suggests that "there is no human faculty of intuition, as the word is ordinarily meant; that is, knowledge acquired without prior exposure to the concept, without having to go through a learning process," people do have a great number of conceptual classifications or models in common. Physical classification schemes can be effectively used in visual presentations, so long as they rely on the same contextual expectations. For example, viewers rely on cognitive intuition to assume that when bar *A* is taller than bar *B* in data visualizations, *A* represents something that is larger than *B*. Likewise, by drawing on cultural cues, such as the association of red with stop, users have a more intuitive understanding of what they are viewing.

Visual cues are given through the physical world to support legibility and readability. Through noticing patterns, people learn to make sense of the world around them. By relying on the assumptions learned in the physical world, and relating them to descriptions of a digital community, we should be able to build an environment that is more "intuitive." This approach is not without faults. For example, the tradeoff of drawing on cultural norms for intuitive representations is that these images are partially subjective. As a result, some of our work presents visual language barriers that are not cross-cultural, and thus require some viewers to actively translate the information from its visual elements.

4.2. Multi-scaled information visualization

To comprehend the quantity of information in Usenet, viewers need to be able to view the space at multiple scales, from broad views of the overall social landscape, down to the level of individual interaction. In order to make these multi-scale views legible, we consider the image of how a city is presented, a metaphor that is commonly used in presenting multiple levels of dense information [23].

Imagine looking at a city from above. From the airplane, the first layers of a city become visible – the density, the relationship between the city and the surrounding areas. As you get closer, you begin to differentiate between various buildings, recognize the streets that separate the buildings, see the forests and the level of traffic congestion. At each level, details emerge

within the larger patterns, allowing you to see the trees that build the forest.

Observing the social patterns of the individuals, how they act collectively and with whom they congregate gives another perspective of insight into a city. Before you see individual people's characteristics, you get a sense of collective attributes. For example, a buzzing street filled with people in suits gives you a collective sense of the environment; looking closer allows you to view individual people's outfits. At that level, you are able to recognize someone who deviates from the collective social norm. The legibility of a space is not only determined by its architecture, but also by the multiple layers of social interaction.

Like its physical counterpart, portraying the online social landscape through multiple layers of interrelated information provides users with an opportunity to see both the overall social character, as well as the more detailed attributes of individuals and groups within the community. All of these levels are important when trying to get a sense of a social atmosphere.

Shneiderman views this multi-layered perspective as essential for information visualizations, emphasizing it regularly as part of his "Visual Information-Seeking Mantra" – "overview first, zoom and filter, then details-on-demand" [19]. Following this advice allows us to build a visualization where users can interact with the space, allowing them to be more immersed in the environment, and thus develop a deeper understanding of its attributes.

4.3. Kinetic typography

In its current state, Usenet's social information consists entirely of textual data. Although users can build abstract mental models of the individuals, the conversations, and the social anatomy of the space, text is still a core component. For this reason, we believe that it is crucial to use typography rather than creating abstract representations.

In his textual landscapes, Small [20] uses the text as his core visual form. By manipulating visual aspects of the text, including the font, color, size, and location, he is able to create an interactive information landscape that allows users to navigate large bodies of textual data.

One of Small's collaborators, Ishizaki, expanded this research to build a design technique known as kinetic typography. Kinetic typography uses the form and function of type and motion to create an expressive visual representation of text. In describing the issues in kinetic typography, Ishizaki's group at Carnegie Mellon constructs a foundation of questions that can be used to start adding expression onto text [10]. Similarly, Rosenberger [16] developed Prosodic Fonts to give textual glyphs visual emotion based on interpretations of a sound signal. Given the textual nature of Usenet, these different research projects help us think about the power of using text for more than its functional attributes.

5. Our design iterations

In the previous sections, we conveyed the big picture – our goal and the underlying philosophies and questions that direct this work. In order to start building such a representation, we must first develop a visual language for representing the social characteristics of Usenet, thinking about how shape, form, color, and motion can be used to express these qualities. Although these pieces show the early stages of ongoing work, we believe that the process and considerations used to build them help set the stage for future work in this area.

In this section, our goal is to describe how we have started to relate social characteristics to visual qualities. At each stage, we focus on the approach and reasoning behind each design, discussing its strengths and weaknesses, and considering how we could integrate this element into a larger visualization. Although there is no final integration, we believe that this work still helps address the fundamental questions we have laid out. The value of this approach lies in the considerations used to build these designs, and in constructing a framework for more extensive research in this area.

5.1. Representing an individual

As we are focused on a people-centric approach, the individual is the core of our visualization. Considering the previously discussed questions for considering an individual's social attributes, what type of form adequately represents the information about an individual's history and interactions? How can motion be added to that form to help convey tone or emotion?

Figure 4. Representing an individual's evolution over time through shapes

In Figure 4, we considered how one's form might evolve over time to represent one's history of participation. We wanted to use the form to evoke strong perceptions about the person's message tone, and thus the person. In developing this image set, we thought about an individual whose initial participation was rather innocuous, but eventually became provocative of insulting through the use of harsh language or tones. Thus, the shape on the far right appears more visually rigid, edgy and sharp.

Initially, we realized that the element on the far right might appear to portray a negative judgment of the

individual. This is problematic because some newsgroups encourage and reward edgy responses, such as the group alt.flame. One the other hand, this also assumes that the visual element has a negative connotation. Perhaps, the image only conveys differences in the tone of users, with the connotation being developed by the viewer.

Developing this piece allowed us to reflect on how one develops and maintains a form, considering how one's frequency and history should be factored in. Should we differentiate a new user with only one edgy message from a frequent poster who is often edgy? What if that active person was only edgy early on and has posted no edgy messages this year?

In order to think about weighting and the magnitude of one's involvement, we started to consider the role of size. With an individual, we could imagine size to mean either length of involvement or intensity. One downside to this approach would be that a person who takes up a larger portion of the screen may appear more valuable – an effect that could be counter-productive if the most frustrating member of the group is the loudest.

One approach is to imagine one's representation as a conglomeration of smaller forms, thus exploring the meaning of density. We imagined building a user's form from the collection of shared words, allowing for density to relate to verbosity (Figure 5). We imagined that the collective form should relate to the individual, potentially constructing one's name.

Figure 5. Using type density to build form

Using text to represent people makes a lot of sense. The tendency to use abstract representations in computational pieces stems from the desire to develop visualizations that can represent generic data sets. Alternatively, hand drawn visualizations frequently relate the visual object to the information being conveyed, so as to help the user understand the representation [22]. Since our work is focused on representing an individual known through word-based messages, it's only logical to use the features that we know about the medium – specifically, text.

Although building one's image through the history of one's posts appears natural, doing so fails to acknowledge the timing of those words. For example, if someone posted hundreds of messages 10 years ago, should their representation be as dense as an active participant?

By considering history, we started to wonder what a decrepit, or dying, form might look like. To express this, we started imagining ways to personify a visual shape, specifically a textual one. Although previous work in kinetic typography never altered the form of the text for functional purposes, we are using text primarily for its form; thus, it makes sense to build personality into the actual glyphs.

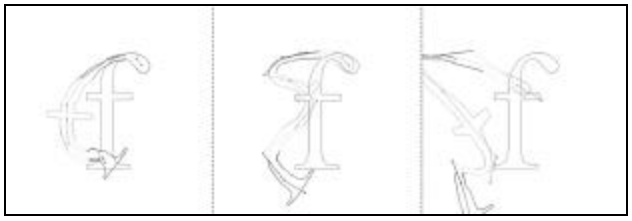

Figure 6. Creating expressive typography through shape manipulation

To personify a textual form, we built a tool where a glyph is given physical attributes, such as joints, flexibility and weight (Figure 6).. The glyph deforms depending on these attributes, and the associated physics equations. Our goal was to base the motion on Disney's [9] "fundamental principles of animation," emphasizing squash and stretch and exaggeration. We focused on how a textual object might move when forced to interact with other textual objects, such as might occur if they were in conversation. For example, imagine an individual who consistently defers to others. Using a rubbery fluid form, this person would appear to ooze and lose form (or identity) when s/he comes into contact with a persuasive individual.

The expressive quality of this motion was quite effective and requires further exploration. We believe that applying different animation techniques to graphical forms has the potential to convey a wide range of emotional attributes, potentially personifying non-human visual representations. One future direction would be to correlate physical properties of a representation with different emotional states, extending from the work done at Pixar (e.g., [11]).

By considering how one represents an individual, we started to reflect on how shape, form, density, and motion can be used to convey basic features of representation. In order to expand these ideas, the next step was to start thinking about how communication between different individuals should be represented.

5.2. Representing a conversation

As users communicate with one another, they share textual content back and forth. In many cases, a post is in response to another post, developing a dialogue between

many individuals. Although we considered how to use one's historical text as an individual representation in the previous section, we also recognize the need to show that messages are not just held by the individual, but shared with the group. Thus, maybe text should flow from one individual to another, or to the group as a whole.

In order to represent the flow of text-based conversations, we started thinking about how text might move, and what paths it might take. Does a message come directly from an individual into the group domain, even when it is a direct response? In Figure 7, we explore the effect of speed, angle, and path on the appearance of flowing text. For example, a message with quick angular motion that gets directed at a handful of users could represent a harsh response to previous posts by those users.

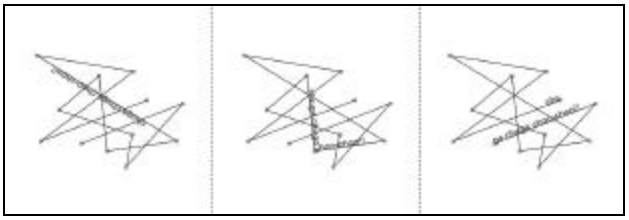

Figure 7. Imagining the path of moving text

While some posts are definitely directed at specific individuals, using a time-based path to convey this quickly becomes convoluted. What order does the path take? How does a message die, or does it constantly just cycle through its path? Here, we started to raise one major problem with using motion in a representation of something that evolves over a long period of time. If we represent everything as though it is happening in the present, messages that are responses to others make little sense. Conversely, if we only represent things in real time, we fail to show the history of interaction or recognize that most people do not read messages in real time. One alternative, explored in Figure 9, is to use an internal clock that moves with some function to time, so that the representation of a week's worth of interaction might occur in 10 minutes of visualization time. We have not even begun to figure out how to represent time well.

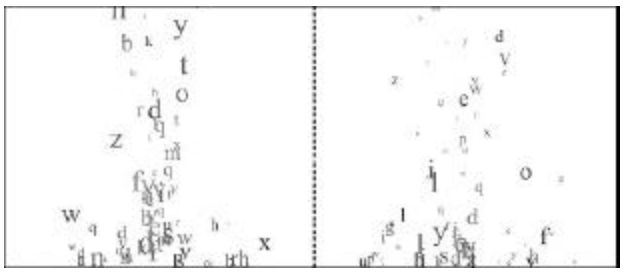

Figure 8. Conversational activity through a fountain of letters.

While still considering time, we imagined a conversation as a fountain of messages, where each new post explodes from the base to become a vibrant and noticeable message, but slowly fades away into obscurity as new posts are written and people stop responding (Figure 8). In doing so, we recognize that socially important messages are not only the most recently posted, but also the ones stirring up active conversations.

To distinguish between recent posts and posts that are part of an active conversation, we built a visualization tool based on actual data (Figure 9). This prompted us to think about how to convey the people who participate in a conversation relative to one another. The idea of a conversation circle evolved from the natural geometry of group conversations. No matter how a population is distributed, if there is something of interest that brings them together, a ring will automatically form, providing the most efficient arrangement for crowds and conversations alike [13]. Like the geometry that forms in physical conversations, each person in our virtual conversation helps construct a larger social circle of people conversing with one another.

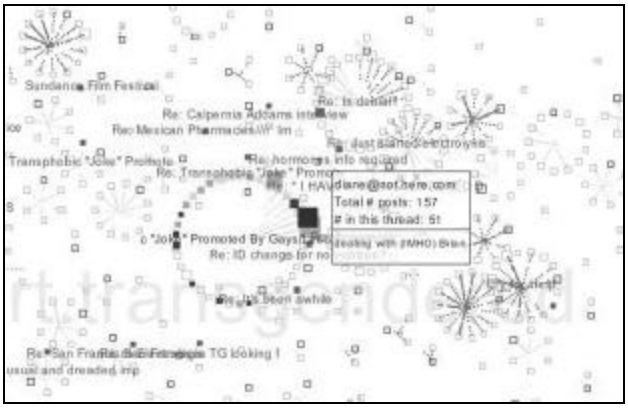

Figure 9. Developing communication clusters

As people converse, the shared information flows from the poster to center of the group, representing a contribution to the common good. We also wanted to emphasize a person's frequent contributions to the conversation, so we related the size of an individual to the number of posts contributed to the conversation.

Although we wanted to represent an individual through dense flexible forms and use actual message text in the flow of conversations, we quickly learned that speed continues to be a limiting factor in large-scale interactive visualizations. Temporarily, we decided to use abstract representations that become text as you narrow your focus on a specific conversation or person.

When we started representing multiple conversations in the same panel, another concern arose. How do you represent an individual's participation in multiple conversations? Physically, it is challenging to be actively

part of multiple conversations, unless they are geographically co-located. Since most people in Usenet are involved with multiple ongoing threads, trying to co-locate them is a graph theory nightmare. Although we temporarily addressed this by splitting individuals into parts, with a physical representation in each conversation and links appearing through interaction, this is neither desired nor intuitive. Coming up with a better way of solving this requires future work.

While using live data, we also wanted to visually explore the statistical aspects of actual newsgroups, so as to compare various aspects of different groups. As a result, we were able to visually differentiate groups based on differences in number of conversations, participants, and messages per thread. Through interaction, we could explore a thread's social network, statistical data on individuals, which conversations an individual participated in, and topics addressed. Although these latter attributes were arbitrarily related to visual components such as color, visible differences between the groups emerged immediately. For example, Figure 9 represents soc.support.transgendered, a support group with long threads engaging many participants, quite a few of whom post multiple times. When we used this same technique to visualize alt.transgendered, a newsgroup overrun with advertisements, no conversation clusters emerged. The entire landscape consisted of individual posts, with only a few two-person exceptions. On inspection, we realized that when people posted more than once, they posted new posts in a series, with identical content in each post. Although these two groups appear to be about similar topics, the social interaction differed enormously.

While we originally built this piece to think about conversations, our curiosity got ahead of us and it quickly evolved into a group representation.

5.3. Representing a group

When considering how to represent a newsgroup, we return to the theories of information visualization discussed earlier. At this level, it is not possible to convey everything that is important within a group. Even if it were graphically possible, we quickly realized that it is not technically possible. Thus, a representation of a newsgroup needs to work on multiple levels.

First, the representation should serve as a fingerprint of the space, conveying the fundamental social aspects of this group so that a quick glance reveals interesting features of the group. This fingerprint is useful when considering how to portray the entire landscape of Usenet, an aggregate of all the different groups.

In addition to using this representation as an overview, it should also serve as the portal to all information within the group – including the specific conversations and individuals. Just as an aerial view of a city portrays the city as the sum of all its parts, so should a newsgroup representation. By building a representation that is a conglomeration of its components, we give the users the ability to see each conversation and person in context, creating an obvious link between the different layers of information. This also encourages interaction, allowing users to explore the finer details of the environment through interactively zooming in on specific aspects that appear interesting, and accessing individual data as desired.

We use dense layers of text to build up the multiple layers of information, even though not all of the text is readable at the highest level. Although it may appear counter-intuitive, Tufte [23] suggests that dense representations make a space far more legible than rough generalizations: "High-density designs also allow viewers to select, to narrate, to recast and personalize data for their own uses. Thus control of information is given over the viewers."

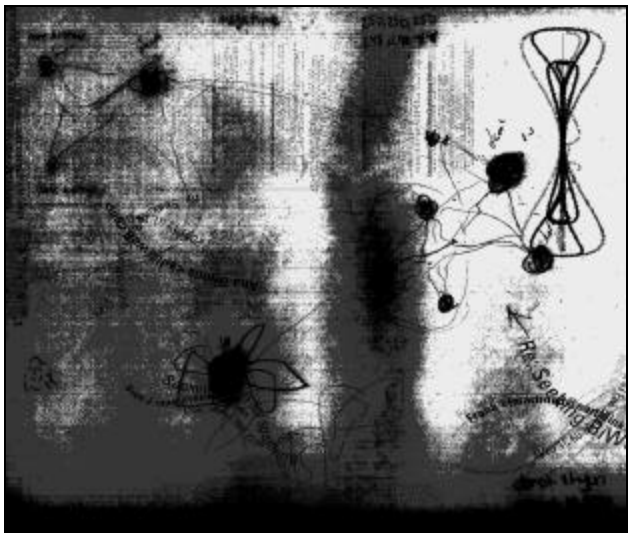

Figure 10. Dense textual information used to construct varied forms

Using data from a newsgroup composed of spam, advertisements and little conversation (i.e. few posts with responses), we start to explore how textual density can be used to build a variety of different forms (Figure 10). Together with the loud colors, these dense forms create an environment that is both noisy and inconsistent, showing off different possibilities while also conveying the general character of the group.

As we thought about density, we considered the historical nature of messages. One of the reasons that a space is cluttered with information is because it is difficult to convey the *now* of an asynchronous communication channel. At the same time, one of the weaknesses of these spaces is that they fail to convey the

history of interaction, making it difficult for users to imagine that a now-dead group was once quite vibrant. Unlike the physical world, interaction in digital environments does not alter the space; scuffs and stains aren't left behind to show use.

Modeled after a dance hall environment, we designed Figure 11 to explore how historical conversations can "scuff" the space by slightly and inconsistently altering the background color over the location of a past conversation. At the same time, we explored how users could affect not only the space, but also each other. Using evolving colors, we imagined that the color of one's representation is altered through interaction, particularly considering what would happen when persuasive individuals interaction with impressionable ones.

Figure 11. Individuals' interactions alter both the space and each other's representation

Although these environment-specific pieces allowed us to explore qualities of that space and of the forms involved, they certainly represent our biases and assumptions regarding the type of social environment. The loud colors and rough forms of Figure 10 convey a very negative quality that we have ascribed to spam-specific environments, a projection that may not be accurate from everyone's perspective. On the contrary, while the computational design shown in Figure 9 showed some group differences, the common elements of the images made it difficult to notice anything but the most extreme differences.

Much work needs to be done at the group level to adequately convey the fundamental differences of the various newsgroups. While using individuals and conversations as building blocks, it is also necessary to figure out how to use layout, coloring and overall style to differentiate between the spaces.

6. Discussion and conclusion

The result of the design iterations presented here is not a final visualization, but a collection of visual ideas that need to evolve and be integrated into one coherent piece, an obvious direction for future work. Alongside these designs, we have articulated important social features for describing Usenet, as well as fundamental ideas that we believe should be applied to future social visualizations. The intention of our method is to start attacking various gaps that need to be filled before we can go from a set of textual data to a visualization that represents the complexity of social interaction behind those words. Through this exploration, we have opened more questions and directions for future consideration, four of which we briefly address in this section.

Numbers do not tell the whole story. While knowing the number of messages in a group may indicate the level of activity, it doesn't help you know if a newsgroup is actually thriving. To know this requires knowing other data, as well as context. For example, thread depth conveys more about a group's success in a support group than it does in a question/answer group. Because of this, visualizations intended to convey underlying social attributes must present more than just numerical comparisons. Likewise, determining what the quantifiable data represents requires much more analysis.

Text is crucial; its form must be explored. While using visual abstractions can help portray features of Usenet, eliminating the text removes the most valuable aspect in conveying social information. Instead of limiting ourselves to thinking about the functional aspects of text, we should think about it as a fundamental form in portraying conversations. Also, by using text to build other forms, we can more accurately convey multiple layers of information, allowing the zooming process to truly present more detailed information. To most effectively use text as form, we must further investigate its qualitative attributes, particularly in relation to font, color and motion.

Personification through motion. By animating a graphical object, affecting not only its position but also manipulating its internal form, we can use motion to grab attention as well as produce a range of expressions. This offers the possibility of personifying text, or visually conveying the personality and tone that the underlying language expresses. If we could visually convey the tone of each message, graphically clustering messages would help evolve the social character of each conversation and group, creating a cohesive visualization full of life. To do this requires applying animation techniques to typography as well as furthering linguistic emotion detection techniques.

Mapping social features to visual characteristics. Although we have started to relate features of social interaction to visual attributes, there is still much more to pursue here. For example, given the use of motion, how should we treat time in the visualization? Should it relate

to the time in which messages were posted or smoothly represent the flow of ideas within a thread? If threads are not related, do they need to follow one universal timing? Other areas that need extensive consideration include color and layout. Can color be used to convey meaning or is it too arbitrary? How should layout be used to convey time, importance, relationships or other features?

We hope that by discussing our approach, and the initial progress that we have made, we are able to motivate others to consider these issues and help us develop better visualizations to convey socially salient characteristics of online social environments.

For further information on the evolution and continuation of this project, please view our website at: http://smg.media.mit.edu/projects/loom2/

7. Acknowledgements

We would like to thank the MIT Media Lab, its sponsors, and particularly the Sociable Media Group for their support, dedication, encouragement, and late night marathons. Special thanks and appreciation go out to Scott Golder, Ming-En Cho, and Ronen Barzel for their technological and editing support that make this project and paper possible.

8. References

[1] Bentley, I. et. al. 1985. *Responsive Environments: A Manual for Designers*. London: Butterworth- Heinemann.

[2] Calvino, I. (1972). *Invisible cities*. London: Picador.

[3] Card, S., Robertson, G., and Mackinlay, J. 1991. "The Information Visualizer: An information workspace," *Proceedings of CHI 1991*, New Orleans, LA, April 27 - May 2.

[4] Chi, E. 2000. "A Taxonomy of Visualization Techniques Using the Data State Reference Model," *Proceedings of InfoVis 2000*, Salt Lake City, UT, October 9-10.

[5] Cooper, Muriel, et.al. 1994. "Information Landscapes," MIT Technical Note. April.

[6] Donath, J., Karahalios, K., and Viegas, F. 1999. "Visualizing conversations," Proceedings of HICSS-32, Persistent Conversations Track. Maui, Hawaii, January 5-8.

[7] Donath, J. 1995. "Visual Who," Proceedings of ACM Multimedia '95. San Francisco, CA, November 5-9.

[8] Erickson, T. & Kellogg, W. 2000. "Social translucence: an approach to designing systems that support social processes," *ACM Transactions on Computer-Human Interaction*, Volume 7, Issue 1, pp 59-83. New York: ACM Press.

[9] Frank, T. & Johnston, O. 1981. *Disney Animation – The Illusion of Life*, New York: Abbeville Press.

[10] Ishizaki, S. 2001. "Issues in Kinetic Typography." Pittsburgh, PA: Carnegie Mellon School of Design. http://www.cmu.edu/cfa/design/kdg/kt/kt_issues.html.

[11] Lasseter, J. 1987. "Principles of traditional animation applied to 3D computer animation," *Proceedings of the 14th annual conference on Computer graphics.* Anaheim, CA, July 27 - 31.

[12] Lynch, K. 1960. *The Image of the City*. Cambridge, MA: MIT Press.

[13] Milgram, S. 1977. *The Individual in a Social World: Essays and Experiments*. New York: McGraw-Hill.

[14] Raskin, J. 2000. The Humane Interface: New Directions for Designing Interactive Systems. Reading, MA: Addison-Wesley.

[15] Richardson, S., Dolan, W., & Vanderwende, L. 1998. "MindNet: acquiring and structuring semantic information from text," *Proceedings of ACL-Coling*. Montreal, August 15-16.

[16] Rosenberger, T. 1998. "Prosodic Font: Between the Spoken and the Written." Massachusetts Institute of Technology: MAS Thesis.

[17] Roth, S., et. al. 1996. "Visage: A user interface environment for exploring information," *Proceedings of Information Visualization*. San Francisco, October 27 - November 1.

[18] Sack, W. 2000. "Discourse Diagrams: Interface Design for Very Large Scale Conversations," *Proceedings of the HICSS-33, Persistent Conversations Track*. Maui, Hawaii, January.

[19] Shneiderman, B. 1987. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, Addison-Wesley Publ. Co., Reading, MA.

[20] Small, D. 1996. "Navigating Large Bodies of Text," *IBM Systems Journal*, Vol. 35, No. 3&4.

[21] Smith, M. and Fiore, A. 2001. "Visualization components for persistent conversations," *Proceedings of CHI 2001*. Seattle, WA, March 31 - April 5.

[22] Spertus, E. 1997. "Smokey: Automatic recognition of hostile messages," *Proceedings of Innovative Applications of Artificial Intelligence (IAAI)*. Providence, RI, July 27-31.

[23] Tufte, E. 1990. *Envisioning Information*. Chesire, CT: Graphics Press.

[24] Wasserman, S., & Faust, K. 1994. *Social Network Analysis: Methods and Applications*. New York: Cambridge University Press.

[25] Whittaker, S., Terveen, L., Hill, W., and Cherny, L. 1998. "The dynamics of mass communication," *The Proceedings of CSCW '98*.