


THE SOCIAL MACHINE



Designs for Living Online

Judith Donath

© 2014 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu.

This book was set in Minion and Scala Sans by the MIT Press. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Donath, Judith, 1962–.

The social machine : designs for living online / Judith Donath.

pages cm

Includes bibliographical references and index.

ISBN 978-0-262-02701-4 (hardcover : alk. paper)

1. Online social networks. 2. Social media. I. Title.

HM742.D56 2014

006.7'54—dc23

2013035294

10 9 8 7 6 5 4 3 2 1

VISIBLE CONVERSATIONS: SEEING MEANING BEYOND WORDS

People participate in many social activities online: they post pictures and videos, write long, impassioned articles, and battle each other in multiplayer games. The predominant social activity, however, is conversation in its many forms: email; chat; discussion postings; status updates; comments exchanged about photos, games, and articles; and so on. Though the form may vary, all share the fundamental quality of being exchanges of words among two or more people.¹ Conversation is much more than a means for exchanging information. It is a form of social grooming: it is how we form and maintain ties, hold our community together, and establish social norms (Dunbar 1996).²

Like face-to-face speech, online conversation is an exchange of words among people. Yet there are significant differences. Most spoken conversations are ephemeral: the words disappear into the past as soon as they are said. Many online conversations are persistent, their text permanently preserved in an archive. Spoken conversations occur in real time: you hear what I say as I say it and can respond immediately. In contrast, many online conversations are asynchronous: I write a message, and you might not read and respond to it until hours, days, or more have passed. Asynchrony adds convenience, for I can maintain a long discussion with you, even if we are never free to meet at the same time, but it also drastically changes the rhythm of discourse.

In this chapter, we will explore visualizing the social patterns in online conversations. These visualizations are the “medium shot,” showing the participants in their situational context and highlighting their relationships and reactions. Our focus will be on visualizing text conversation, the ubiquitous medium of online comments, email, discussion boards, and so on. We will

also look at how visualization can supplement even face-to-face discussion, bringing our attention to the nuances and behaviors of which we are often unaware.

TEXT AS CONVERSATIONAL MEDIUM

The primary medium of online interaction is text. Text conversations are interactions stripped of almost everything but the words, a simplicity that yields many advantages. They are easy to implement; pioneering programmers developed the earliest systems in the mid-1960s to send messages among users of time-sharing mainframe computers (Van Vleck 2001). They are also easy to participate in, requiring only a keyboard. The sparseness of text conversations focuses attention on the words themselves, without the distraction of the speaker's appearance. No one can dominate the conversation simply because he or she is bigger and has a louder, deeper voice. In discussions where the participants are strangers to each other, text interactions make it possible to converse without being swayed by knowing the others' age, race, or gender.

Simple as they are, text conversations are an example of a design that goes "beyond being there." The writer of a message can carefully edit it before sending. Recipients can read a complex text slowly and repeatedly, yet quickly skim routine communications or dull sections. They can save messages, forward them, or send them to a computer for analysis and visualization. Text conversations are frequently asynchronous, so one can read and write at one's convenience. Text conversations can be dialogues between two people or a mass discussion involving hundreds or more. Text is flexible. You can use it to write a formal document or to toss off a slang-encoded message. It's easy to quote someone else's words to respond to a particular point. Text requires only a keyboard to produce. Reading text is fast; experienced readers can skim paragraphs quickly.

Yet the sparseness of text conversations has disadvantages, too. Face-to-face conversations convey subtle but important social information nonverbally, such as the speaker's emotions and the participants' relationship with each other; these are hard to discern through text alone. In person, we gather much identity information simply by seeing each other—even when talking to a stranger, you can infer gender, age, race, and the myriad cultural affiliations we advertise with our clothes, hair, and the like. When strangers

converse via online text, they remain ciphers, lacking the identity information that provides essential context for understanding each other's meaning.

Text conversations have little of the rich nonverbal communication that face-to-face conversation offers. Gestures, facial expressions, gaze, accent, and tone of voice all contribute to the social meaning of an exchange. We look at other people's facial expression to judge their sincerity, to assess whether we are holding their attention (and if not, to see what has drawn their gaze). We add meaning to our words with tone of voice; we can make a "thank you" heartfelt, perfunctory, or ironic, simply by changing the way in which we say it. Furthermore, these numerous and subtle forms of nonverbal communication exist within a set of cultural norms for their performance: we assess others in terms of how they conform to expected behavior. Gaze alone has an enormous number of rules: how long do you hold someone's eye when you speak to him, when is looking away "shifty," when is looking at someone paying attention, and when is it rudely staring? Where on a person may you look? Deviating from the norm may make the impression that one is aggressive, shy, lying, and so on. A vast amount of subtle interpersonal communication occurs outside the realm of the words themselves and is simply missing from the abstract space of text-based communication.

VISUALIZING TEXT CONVERSATIONS

Visualizing text conversations can bring richness and subtlety to this medium by making visible the existing but hard-to-discern social patterns within them.

Although social cues are sparse in individual interactions, many social patterns exist in the history of the interactions. Visualizations compress this history, making it immediately accessible in the context of the ongoing conversation. We can design visualizations that trace the patterns in, for example, individual participants' actions, the changing topics of conversation, and the temporal rhythms. The challenge is both to identify which are the meaningful patterns and to represent them in a form that is both intuitively readable and yields a more nuanced understanding of the conversation's social dynamics.

Visualizing text can help convey the emotions behind the words or the ebb and flow of excitement in the discussion. The overall appearance of a screen full of comments angrily arguing about mandatory vaccinations or Israeli settlements is not very different from one full of readers' funny pet

FIGURE 6.1

Bernard Kerr, *ThreadArcs* (2003). *ThreadArcs* is a compact and graceful depiction of an email conversation's threading structure. Each dot is a message and the arcs show replies. One can quickly see the characteristics of each conversation—in some, discussion focuses around a single message, whereas others consist of independent comments with occasional responses (Kerr 2003).



stories. By highlighting the rhythms and emotions in a discussion, visualizations create a vivid and legible social space.

Participants in online discussions come and go, making it hard for a newcomer to grasp the relationships and dynamics within the group. And even regular participants may have trouble picking up the local social mores because subtle expressions of encouragement or disapproval are hard to see. Visualizations that show the participants, their roles, and the patterns of their interactions can make the social information that is embedded in online discussions easier to perceive.³

The archives of a conversation are a rich source of information, not only for finding specific facts (which search interfaces enable), but for seeing the development of a decision and the structure of the community. But wading through this material can be tedious. An interface that makes the archives more interesting and that highlights key information encourages people to read and understand the context before joining in.

Here we will examine three examples of conversation visualization. *Newsgroup Crowds* depicts the tenor of a discussion space by analyzing the basic social roles its participants play. It is designed to help participants more easily see the social groupings and mores of a large-scale conversation. *The-mail* analyzes email exchanges to create a portrait of a relationship; it makes it easy to explore one's personal archives. Finally, *History Flow* depicts the evolution of a collaborative production, revealing the controversies and social dynamics that lead to the final text.⁴

These examples were made as stand-alone visualizations, but it would be easy to integrate them into a live conversation space. One could display *History Flow*, which shows changes in Wikipedia articles, at the top of an article's history page; *Themail* could be part of an email interface. As supplements, the role of the visualization is to add context, clarity, and vividness, without taking away from the simplicity and versatility of the medium.⁵

Newsgroup Crowds

From the late 1970s through the 1990s, Usenet newsgroups were the most popular form of online conversation (Hauben and Hauben 1997; Whittaker et al. 1998). These large-scale, public discussions covered an immense range of topics, from technical discussions of computer systems to political arguments and child-rearing advice. Although all the groups shared a common technology—plain text threaded discussion—the tone and interaction patterns varied greatly between groups. Some were lively, sociable discussions that flowed freely from topic to topic. Others were more cut-and-dried question-and-answer forums, where people went for advice on a specific problem. Some had been taken over by “flame wars,” in which a few participants argued furiously and endlessly. Some were vibrant and populous, with numerous messages appearing every hour, while others were withering away.

As is typical with many large discussion sites, if you were seeking something specific—if you had a question about writing compilers for the computer language C or verb forms in Serbian—you could probably find the right newsgroup to ask. But finding a congenial community—a place you could return to day after day, not just to ask questions but also to answer them, and to enjoy the banter among people who shared some interest or outlook with you—could be harder.

The *Newsgroup Crowds* visualization depicts the interactions within individual groups, making it possible to see their ambience at a glance (Turner et al. 2005; Viégas and Smith 2004). It is a simple bubble chart, a form of statistical graph often used for financial projections and other business statistics. But in this case, it creates an evocative view of the discussions. Its strength is in its careful choice of what data to depict.

Each circle represents an individual. Size shows the overall number of posts the person made, and brightness shows how recently he or she participated. The visualization's expressive power is in the circles' placement.

FIGURE 6.2
Fernanda Viégas and Marc Smith,
Newsgroup Crowds (2004). A popular and
lively group, with both deeply engaged
participants and occasional visitors
(alt.politics.bush).

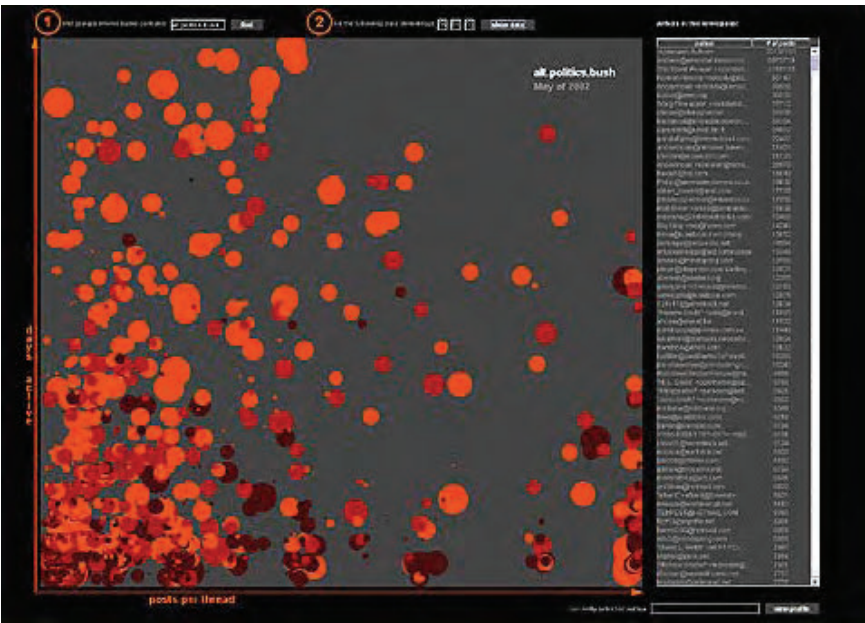
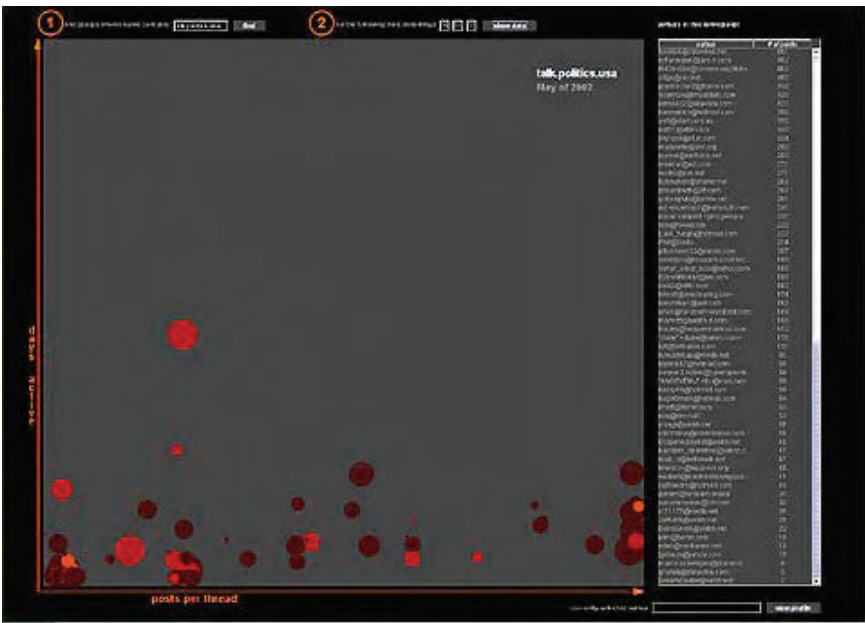


FIGURE 6.3
This group seems moribund, with no
recent posts and mostly short-term visitors
(talk.politics.usa).



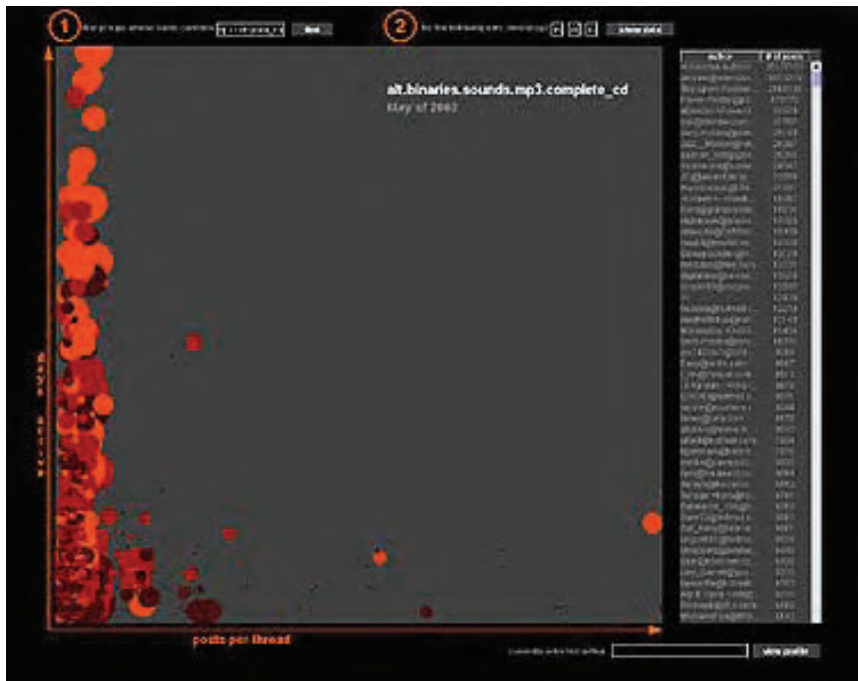


FIGURE 6.4

The clustering of all participants to the right shows that this is not a conversational group: people post only once in a thread. And, in fact, it is a binaries group, a place to upload media, not converse (alt.binaries.sounds.mp3.complete_cd).

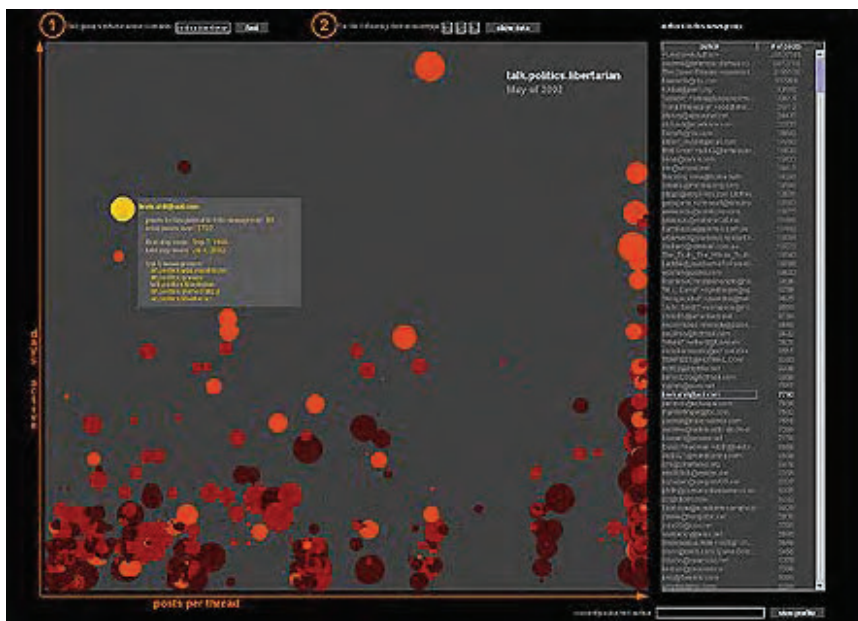


FIGURE 6.5

A group dominated by an inner circle of very active participants—quite likely an argumentative group. Other people do participate, but more occasionally (talk.politics.libertarian).

Location on the y-axis shows the number of days the person has been active in the group over a chosen time period. The x-axis location shows the average number of posts he or she contributes per thread, a statistic that provides an interesting measure of conversational participation: the amount of back-and-forth discussion in which the person engages.⁶

The *Newsgroup Crowds* visualization forms clusters of participants based on their characteristic behaviors. Occasional visitors, those who come to ask a question and then are gone, are in the lower left. People who answer questions are typically in the upper left; they return habitually, but contribute few posts per discussion. Engaged participants are further toward the right, with more messages per thread, for they participate in back-and-forth conversational exchanges. However, among the most engaged are the flammers, people who argue viciously and destructively. Nonconversational—and indeed, not necessarily human—contributions appear at the extremes near either axis: these are postings of binaries,⁷ newsfeeds, and other automatic scripts.

Understanding the social roles the participants play in a discussion helps you know what to expect from them and provides context for interpreting their words and behavior. Seeing the distribution of roles within a group—leaders, supportive people, cranks—can help you understand its social dynamics. A group with two or three strongly opinionated leaders will have lively, but possibly constructive, disagreements. One with too many cranks and other irritating participants is likely to quickly lose members.

The distribution of these different types creates a group's ambiance. As with diet, variety is best. A diverse set of participants—some devoted, engaged participants, some quieter but pithy answerers, and newcomers for fresh ideas—makes for a vibrant conversational setting (figure 6.2). One where few people are posting and not much has been added recently is probably dying off (figure 6.3). One that has many postings but few ongoing threads is not a social space; this is the signature of a group that functions mainly as a repository for binaries (figure 6.4). A group dominated by a few debaters may be an interesting discussion between a few experts or, more likely on Usenet, a flame war (figure 6.5).

This visualization is simple but clear. It is easy to grasp that a chart with a few faded bubbles is moribund, and that a crowded, bright one is lively. Viewers need to learn the meaning of the axes, but once they understand that, the depictions of different conversation environments are legible and expressive.

both with the subject's other correspondents and is common in general English usage (see figure 6.6).⁹ The columns of yellow words show the words that typify a month's correspondence. The height of the columns shows the volume of email exchanged that month, creating a histogram of the frequency of their correspondence. The circles represent individual messages; their color indicates whether they were received, sent, or came via a list, and their size indicates the length of the message. The green words in the background are words that typified the subject and correspondent's exchanges over the course of a year.

A quick glance shows the extent of the relationship: whether correspondence has continued over years or was a brief flurry over the course of a few weeks, whether it is regular or intermittent. The background words give a basic picture of the content of their discussion. But a closer look at the words in the timeline really makes it come to life. These lists sketch the outlines of a personal narrative. Some are banal, such as the many columns dominated by "meeting," "schedule," "time," and "room" in one's correspondence with the departmental secretary. Others recollect travels, with distant cities and foreign money units, or events such as weddings, with words like "invitation," "seating," and "honeymoon."

We use photographic images to record our shared experiences in the physical world. *Themail* visualizes our shared interactions in the online world, an increasingly significant part of our lives.

Users of *Themail* found the portraits recognizable and evocative, mentioning that the depicted histories reminded them of forgotten events or captured incisively what makes a particular relationship special.

The most unexpected thing for me was simply the amazing feeling of launching this visualization and seeing, for instance, the exchanges with [my wife]. There were words like "love," "hope," "marriage," "change ..." It was great! It managed to sum up in a few words a lot of what was being said at that time.¹⁰

Others noted that it helped them see how their relationships evolved:

I was really interested to see this one. During the past five years Ray has gone from being an acquaintance to a very good friend. Looking at it actually takes a while for the words to be dominated by things like bar names, beer and

cinema! There are a couple of things that come out of the visualization, like a holiday when we all went to Sri Lanka and when Ray went to work in another town for a few months.

People usually think of their email as private, and indeed *Themail* was designed to be a personal visualization for perusing one's own email. But as with photographs, many of the users of *Themail* wanted to display the portraits and share them with others. *Themail* skirts the boundary between public and private; by using only single words, it depicts the essence of a conversation without revealing the specific meaning. It creates displayable virtual snapshots that reveal enough of the relationship to be interesting to an outside viewer, but which also maintain the privacy of the subjects' correspondence.

To be honest, I shared stories that I discovered in the archive with family members and with a few colleagues. I was moved to talk about the content of some of the messages, much like someone would be moved to share a memory sparked by an old letter or photograph.

Themail depicts personal history, reminding the user of forgotten events and previously unnoticed patterns. The next visualization, *History Flow*, depicts the public history of Wikipedia edits, showing the social interactions that shape an article's evolution.

History Flow

Wikipedia is a collaboratively written encyclopedia: hundreds of thousands of individuals write, edit, and revise its articles (Suh et al. 2009). It is a fascinating and highly functional model of collective work. There are no barriers to participation: anyone can edit it. Some contributions are poorly written, others are badly researched or express personal opinion rather encyclopedia-style fact; some are deliberately and maliciously wrong. Yet over time, the collective editing process hones the articles to high standards that are arguably equivalent to *Encyclopedia Britannica* (Giles 2005).

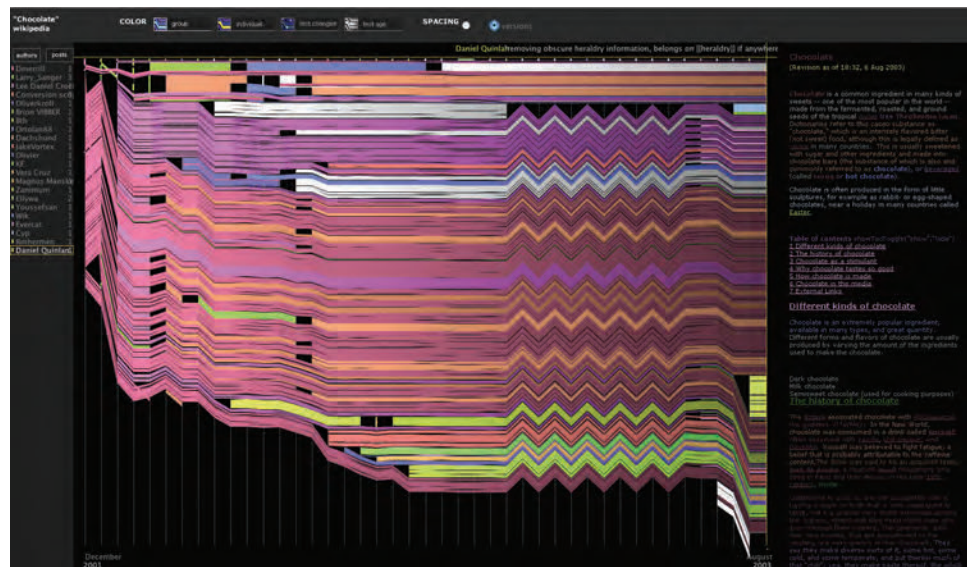
We can think of the articles themselves as discussions in which participants communicate through both additions and deletions. (There are also traditional discussions on the Wikipedia site. Each article has a discussion

page where contributors can explain the reasons behind their edits.) Wikipedia articles include a link to their history of editorial exchanges, but in raw form these are far from compelling reading. A contributor involved in a heated edit war might be motivated to delve into this material, but few ordinary readers would. Yet this “conversational” history is rich with detail about how the article came to be. Is it primarily the work of a single person, or has a broad community created it? Are sections of it controversial? What have the arguments been over? Is the controversy between two or more sincere though conflicting views, or has the article attracted malicious vandals?

History Flow visualizes the edits that have produced a Wikipedia article (Viégas, Wattenberg, and Kushal 2004).¹¹ Each contributor is represented by a color, with anonymous ones shown in white and gray. Vertical slices show the state of the article at a particular time, and text that persists between versions is connected between slices. We can thus easily see whether an article has a few authors or many, and whether they are named or anonymous. We can also see periods of growth. Most interestingly, we can see areas of controversy, where prose has been repeatedly added and deleted, as well as instances of vandalism—and the speed at which they are repaired.

FIGURE 6.7

Fernanda Viégas and Martin Wattenberg, *History Flow*, “Chocolate” (every edit) (2004). This figure and figure 6.8 show the history of edits to the article about chocolate. This version shows the state of the article each time there is a new version. On the right, you can see the characteristic zigzag pattern of “edit wars,” where a change is made and then repeatedly removed.



History Flow illustrates how an article changes over time in two ways. One shows all the edits that have been made (see figure 6.7); when there is an “edit war”—repeated edits and reversions—the first mode displays a characteristic zigzag pattern. The other mode shows a snapshot of the article at equal time intervals and a blank column where a page’s content has been deleted (see figure 6.8). This version presents a much smoother picture of the history. Here, the edit wars are often invisible, since the reversions are made so quickly that we do not see them in these time slices.

This vividly demonstrates how a simple design change can tell a different story; both versions are useful, for each depict Wikipedia from a different perspective. The first version, which shows all revisions, no matter how briefly they were up, tells us about the writing experience: where was there controversy, how many authors have worked on the article, and so on. The second, temporal view tells us about the readers’ experience. If vandalized pages are found and corrected quickly enough, few if any people will see them (Priedhorsky et al. 2007); this view shows which versions have been online long enough to have been read.¹²

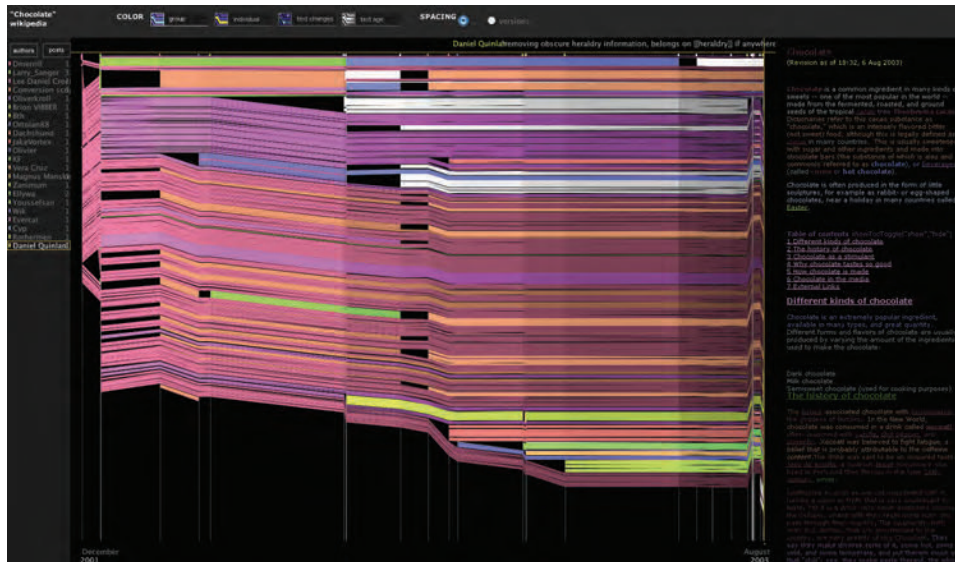


FIGURE 6.8
Fernanda Viégas and Martin Wattenberg, *History Flow*, “Chocolate” (temporal sample) (2004). This version shows the article’s state at measured time intervals. It demonstrates that viewers are unlikely to have noticed the edit war: sampled over time, the article appears far more stable.

Wikipedia's transparency is one of its great benefits. We usually see finished products, whether articles, movies, or car designs, with little idea of the process that created them or the controversial decisions made along the way. Wikipedia's archives let us delve into the making of an article. At first, the idea that one might want to peruse the editing history of an encyclopedia article might seem unlikely, but *History Flow* helps you immediately see where interesting things have taken place, and it's easy to imagine the viewer of this visualization becoming intrigued about what occurred. Making the history of creating an encyclopedia article into an interesting narrative encourages readers to investigate its reliability. With a tool such as *History Flow* they can quickly see what statements are controversial; they can see if an article was written by a couple of people or by hundreds: the highest-quality articles have had extensive editing by many people (Wilkinson and Huberman 2007). The visualization shows the importance of seeing the social process that created the finished article.

VISIBLE SPEECH AND PRESCRIPTIVE DEPICTIONS

The value of visualization with online text discussions is clear: it can make typically hard-to-perceive social patterns visible, thus providing additional social cues to an otherwise sparse medium. But can visualization add anything to the already rich experience of face to face encounters?

The most useful visualizations show us patterns beyond what we typically observe. Try this design exercise: Follow two conversations, one online and the other among people meeting face to face, and sketch them in any way that seems meaningful. When I assigned this to my students, the differences between the two sketches were striking (Donath 2004b). Almost invariably, the sketches that they made of the online conversation focused on the discourse, delineating the ebb and flow of topics. The sketches they made of face-to-face conversations focused on the people, with the individual participants' and their words each occupying a defined space (see figures 6.9 and 6.10).

The sketches reflect what we typically observe. Online, we primarily see streams of words, and the people are not immediately perceived as distinct individuals. This suggests that a key role for conversation visualizations is to portray the participants more vividly, in order to help the viewer see the conversation space as populated by unique individuals. We will look at several approaches to doing this in chapter 8, "Data Portraits."




FIGURE 6.9

Francis Lam, face-to-face conversation sketch (2004). Each colored circle represents a person; it is pink when they are speaking, and the circle grows larger the longer they talk. The white circle moves to show where the person's attention is focused.

FIGURE 6.10

Francis Lam, online conversation sketch (2004). Here the focus is on the flow of words, rather than the people. Each box is a message, its size corresponding to the message length and its color indicating the author. Boxes stacked to the right are supportive notes; those on the left are more critical. It has a subtly prescriptive air, for one wants to make the tower of messages well balanced.

On the other hand, in face-to-face discussions we are typically aware of the other individuals but often have less of a sense of the conversational dynamics, especially our own contributions. Do I speak more or less than most other people? Do I interrupt others, or are people excessively aggressive in interrupting me? Visualizing spoken conversation to show topic flow or interruptions can reveal surprising insights into how a conversation is controlled and the roles different participants play.

One scenario where visualizing speech would be quite useful is listening to archived discussions, such as a missed meeting or class discussion. Here, visualization can provide a useful way to find key moments. As with the text visualizations we have seen, such a depiction can show where there are moments of excitement and controversy or highlight the roles different participants take in the discussion. For spoken discussions, however, we would use different data: rather than edits and response counts, we have the rhythm and volume of voices, the overlap of interruptions, and the like.

Beyond enlivening the task of listening to discussion archives, visualizations can enhance ongoing conversation, both mediated and face-to-face ones. Here the visualization may play a more activist role: by making otherwise hidden patterns visible, it can influence the participants' behavior (e.g., Bachour, Kaplan, and Dillenbourg 2010; Bergstrom and Karahalios 2007; Brandon et al. 2011; DiMicco et al. 2007). Even when we are present in a face-to-face conversation, we can miss key social patterns and cues. Someone may monopolize the conversation without being aware of it; men may be unaware of how frequently they interrupt women, for example. One goal for making these patterns visible is to provide a catalyst for more equitable interaction.

My students and I were once asked to make a device for “the kitchen of the future.” We decided to make a social interface for keeping company with distant friends while cooking or cleaning up. The telephone does this, but we felt that audio alone is not ideal. We wanted to connect the spaces, not just individuals, so that if there were multiple people in one or both spaces, the virtual conversation would include them as well. A speakerphone does this, but a speakerphone left on indefinitely feels like surveillance; it is easy to forget that you are in a connected space in which your words travel beyond the people physically with you. Though videophones have long been a staple of futuristic domestic scenes, we felt they were poorly suited to this scenario. They work best when the users look into the camera, which is impossible when you are chopping vegetables or washing plates. We felt that a visual yet abstract interface would improve the experience in several ways.




FIGURE 6.11

Judith Donath, Karrie Karahalios, and
Fernanda Viégas, *Visiphone* (2000).

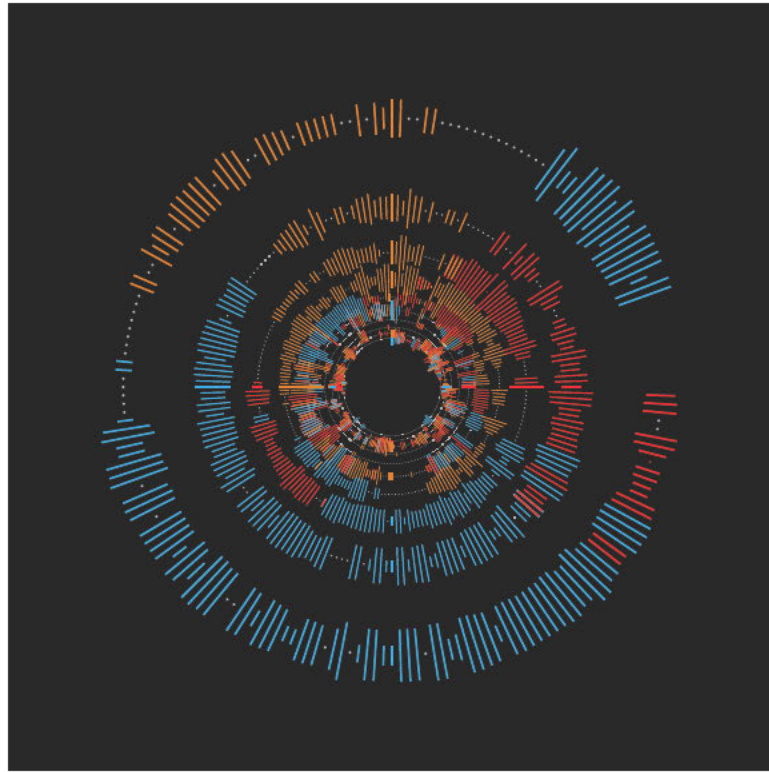
Our solution, the *Visiphone** (see figure 6.11), was a dome with an abstract design that showed the state of the connection (Donath, Karahalios, and Viégas 2000). The circle at the top represents the current audio condition: colored circles indicate that there is sound, with blue meaning local and orange meaning from the other location. The size of the circle represents volume and moments of silence appear as small gray dots. If two people speak at once, the circles are superimposed. The dots spiral down from the top, so the whole spiral is a visualization of the patterns of alternating speech for the past few minutes.

Our goal in designing the *Visiphone* had been to create a visual focus for and reminder of an otherwise invisible audio connection between two spaces. But when we showed this project publicly, many people were especially interested in the overlapping circles, explaining that they welcomed the ability to show overbearing companions how frequently they interrupt.

Karrie Karahalios, one of *Visiphone*'s creators, subsequently created the *Conversation Clock* (Bergstrom and Karahalios 2007), a table-top conversation visualization designed specifically to give participants in a small face-to-face discussion real-time feedback about "turn-taking, domination, interruption, and activity." Each ring is one minute, with a colored tick representing each person who was talking at a given time (see figure 6.12). At the

FIGURE 6.12

Tony Bergstrom and Karrie Karahalios,
Conversation Clock (2007).



end of the minute, the circle shrinks toward the center and a new circle begins. It is easy to see at a glance who dominates the conversation, who interrupts, where there is excitement, and when there is silence.

Although the design is intended to be prescriptive and to influence people to adjust their behavior, it does not explicitly say what the ideal behavior is. The participants must already have the knowledge that, for example, interrupting is impolite.¹³ The visualization simply makes the existing behavior patterns visible and clear; it is aimed particularly at improving self-awareness: I may know not to interrupt, but be unaware that I do so frequently, or that I talk much longer than others do. With no explicit labels of how one ought to behave, it allows for people with different roles or personalities to determine their own participation goals. Someone who is junior or who prefers minimal, concise contributions may be quite content to speak less than others do.

Other designs make their notion of desirable behavior clear. *Second Messenger* created histograms of spoken participation and labeled levels as over- or underparticipating, with the goal of equalizing participation (DiMicco et al. 2007; DiMicco, Pandolfo, and Bender 2004). The results were mixed: those who spoke more than most without the visualization reduced their contribution, but in some cases those who spoke little without it also spoke little with it. Sometimes the more equitable distribution of speaking time was beneficial for the group, but there were also examples where a group's ability to achieve the experimental task was better without feedback about their discourse style.

Awareness applications do influence our behavior, because they provide a guide where there was not one before. ... Providing any automated awareness feedback to groups will cause them to adapt their behavior to accommodate to the normative pressure implied by the feedback. (DiMicco et al. 2007)

Such results are vividly apparent with visualizations of live speech, but the same benefits and caveats hold for visualizations of written conversation, when the visualization is integrated with the user interface.

Visualizations can emphasize or diminish authority, welcome or discourage contributions, and otherwise influence behavior. The choice of what interface to use—or indeed, whether to meet face to face or online—requires first knowing what social structure and interactions one desires.

In face to face meetings, high status members dominate the conversation. Lower-status members participate less and are likely to offer opinions gauged to agree with those higher up, rather than presenting their own assessment. Online, the power of these hierarchies is flattened (Dubrovsky, Kiesler, and Sethna 1991; Sproull and Kiesler 1991a,b). Text discussions generally see greater participation from members at lower ranks—at least in part because the visible reminders of status and identity are greatly reduced.

Is this desirable? It depends. If the lower-status members have useful contributions and critiques to make, then yes, the more open participation of an online discussion can be very valuable. And certainly if the hierarchy is based on unfair or irrelevant properties, flattening it is beneficial (though those at the top are unlikely to see flattening the existing structure as a better order of things).

However, hierarchies sometimes reflect expertise differences that are useful to preserve: when those at the top are the most knowledgeable, the

discussion may work best with contributions coming mostly from the higher-status members while the others primarily listen and learn. Having many people chiming in who know little about a problem may be far less productive than having the most authoritative voices dominate.

Furthermore, getting input may not be the point of the discussion. The goal of some meetings is to convince the attendees to all buy into something; though some may have reservations, for better or worse the goal of the gathering is not to air them.

If you want to elicit a broad range of ideas from lots of people, an online forum may be more productive than a face-to-face meeting, where people are more likely to defer to the ideas expressed by more senior individual. But if you want to get a group of people excited about something—defects and all—then an in-person meeting, where leadership is more powerful, is likely to be more successful.¹⁴

A social visualization used in the context of interaction will affect the participants' behavior. It will do so if it advocates a particular behavior—but also if it simply reflects the group's behavior back to it. Exactly how individuals will react depends on the ideals explicitly or implicitly advocated by the display, the group's norms, and their own personality and role in the group.

TOWARD MORE INSIGHTFUL ANALYSIS

A social interaction generates an enormous amount of information that one could potentially visualize. Some is useless for social purposes, such as counting the distribution of letters in the words that people used (though a linguist or cryptographer might find something to study here). Some is socially relevant and easy to measure, such as how many postings each participant contributed. Some is relevant, but is much harder to assess computationally, such as who is supportive or has the best sense of humor.

The visualizations we have discussed so far have depicted activity and word-use patterns, but have not attempted to analyze meaning or sentiment. As participants in a conversation, we try to understand what is going on—what are people talking about? Are they happy about something? Critical? Is a particular person warm and supportive or frequently sarcastic? Answering these questions, particularly in the sparse environment of a text conversation, can be difficult even for a human reader.



FIGURE 6.13

Alex Dragulescu, *Spam Plants: Untitled I* (2006). Good designs use form and metaphor to convey the meaning of the data, making an image that gives the viewer an intuitive feel for its significance. But visualizations are inherently arbitrary: one can pick any data to show and render it in any form. Alex Dragulescu subverts the usual goal of visualization to create beautiful flowers from statistical analyses of spam email and computer viruses.





FIGURE 6.14

Alex Dragulescu, *Spam Plants: Untitled VII* (2006).

FIGURE 6.15

Alan Mislove, Sune Lehmann, Yong-Yeol Ahn, Jukka-Pekka Onnela, and J. Niels Rosenquist, *Pulse of the Nation* (2010). These cartograms (maps in which the geographic area is distorted to show some other variable, in this case number of postings), show the mood in different regions in the United States at various times of the day, as inferred by analyzing Twitter postings. Though reliably assessing the mood of any single posting is beyond current computational (and, sometimes, human) ability, given a large collection of postings, it is possible to make a good assessment of overall tone (Mislove et al. 2010).



For a computer, such understanding is a work in progress. Analyzing text to assess social and psychological characteristics and the like is an area of active research.¹⁵ It is especially important to be aware of how accurate (or not) such an analysis is, and not to highlight misleading or erroneous data. For example, we might want to create a visualization that depicts emotion, to help us determine whether a community is upbeat or negative, or whether a writer is in an unusually good mood (see figure 6.15). To do this analysis, one might create a lexicon of positive and negative words, and score writings based on the presence or absence of these words. Yet, as this example from an article on sentiment analysis (Pang and Lee 2008, 22) shows, such counts can be misleading: “This film should be *brilliant*. It sounds like a *great* plot, the actors are *first grade*, and the supporting cast is *good* as well, and Stallone is attempting to deliver a *good* performance. However, it can’t hold up.” Though it has multiple positive words, it is a negative review. And just looking for negation at the end will not help; had it ended, “In fact, it couldn’t be better,” it would be positive. Quantity helps: a large dataset of comments will in the aggregate have more positive words if the overall reaction is positive, though any one comment may, like the example above, be misinterpreted.

Newsgroup Crowds mapped a rough assessment of participants’ roles using the straightforward measures of total postings, postings per thread, total days present, and recency; these data are easy to track and socially

relevant. Their clustering was useful and intuitive. But human readers, closely studying the interactions in the group, would distinguish among a somewhat different and more subtle set of roles.

“Celebrities” in online communities post frequently.¹⁶ However, it is the quality of their writing, and not only its quantity, that makes them central figures. They are adept at the type of communication favored by the group, whether humorous repartee or rigorous scientific reportage. Furthermore, they are famous within the community: people talk about them. These influential participants shape the tone of the group, not only through their high level of participation, but also because their high status motivates others to emulate them.

“Ranters” also post frequently. And people talk about them, too. But these irritating participants have low status within the community. Unlike the celebrity, the ranter’s tone or topics are outside of the group’s norms: the atheist in the Christian worship group, the fundamentalist in the evolutionary biology group, the bottle-feeder in the breast-feeding mother’s group, and the like. The distinction is as much one of tone as of topic: a gracious dissenter could be a valued member of the group, though the eloquence needed to sustain such a position in the volatile world of text discussions makes this rare.

A human reader familiar with the group would be able to identify these two (and many other) roles. The problem for visualization design is that statistically, the high-status celebrity and the low-status ranter look quite similar: both post and are mentioned frequently. The key distinguishing element—the esteem in which they are held by their community—is harder to measure. Ideally, the computational analysis should be able to distinguish, as the human reader can, between respect and rebuke; though not impossible, this is far from trivial.¹⁷

Adding another data element can help. Looking again at *Newsgroup Crowds*, some types might be more distinct if, for example, the visualization also indicated how many responses a user’s posts received. Well-regarded participants often generate more responses, especially in a mature community that has established the norm of ignoring annoying or provocative postings. That said, while this simple statistic makes some additional differentiation, alone it is not an accurate role detector: unwelcome and provocative postings can also generate a flurry of reactions. A good approach is for the visualization to remain neutral about the desirability of having few or many responses; its role is to make the patterns clear, and the human viewer interprets them.¹⁸

VISUALIZATIONS SET THE SCENE

The setting of a conversation influences how the participants feel about each other. In person, we arrange the seating in our living rooms to be comfortable and go to restaurants to talk over cocktails; these settings enhance enjoyment and dispose us to be sociable. The physical setting—classroom, church, beach, elevator—also provides cues about how to behave. At a formal presentation, it is normal for some participants to only listen and not speak, whereas during a social dinner such silent non-interaction would seem strange. We learn how to act in certain settings and apply this knowledge to novel experiences: walking into a new classroom with chairs and desks organized in a familiar way tells us that the behavioral protocols we have learned in similar spaces are likely to apply in the new space.

Online, social visualizations can help establish the setting. Long-term discussions that share a common interface look the same, even though over time they have evolved quite different standards of conduct—ideas about what is funny, on-topic, acceptable, and so on. Visualizations of their individual histories can highlight the distinguishing social patterns.

We learn the mores of a setting by observing others and by noting their reactions to what we say and do. When we are face to face, others' reactions provide cues that tell us that we have been talking too long or using the wrong vocabulary: they start to fidget, look away, or appear displeased. But online, you don't hear any laughter if you amuse everyone and you don't see any raised eyebrows if people are offended, at least not until someone is sufficiently entertained or affronted to respond in writing. The difficulty of making subtle responses is the source of many angry online exchanges: offending behavior continues until it is truly egregious, and then meets with an infuriated response. Both the asynchrony of online conversations and the difficulty of conveying subtle hints via text make it difficult for newcomers to pick up the social mores of an online situation. Visualizations that highlight response patterns can help convey these social cues.

Common ground—the participants' shared beliefs and information—is essential to communication (Clark and Brennan 1991). When people communicate cooperatively (that is, when they are not deceiving each other¹⁹), they increase their common ground. Imagine asking someone on the street for directions. At first, you have little common ground. The initial task is to get the other's attention; you now have the common ground of agreeing to

participate in some interaction. Next, you establish that you have a language in common. If they do not understand your initial query, you may try another language you know or write down the address you're trying to find. Once you've established a common understanding of your question, you then establish whether the stranger can help you. If so, they may say something like "Walk down this street until you get to that big tree, then turn right at the next street." These directions assume you share a common physical setting. You might spend some time working out that the tree you are looking at is indeed the tree the other person is referring to.


Throughout this experience, how you speak refers to and helps establish your common understanding of the social situation. If the two participants are college students on campus, even if they are strangers, they recognize each other as part of a larger community whose members easily provide this sort of information to each other. Though polite, they need not be elaborate; they need not explicitly establish the basic common ground of a cooperative relationship. Strangers in a high-crime area, however, might make more of an effort at the beginning of the interaction to establish their good intentions; for them, creating the common understanding of a harmless interaction takes more work. The common ground of a communication has numerous components: common surroundings (being in the same physical space or sharing a virtual space), cultural beliefs, language, and a growing understanding of each other's needs and abilities.

Cultivating cooperation is a form of common ground; it establishes your shared goals. Though the rituals of politeness may decrease common ground in factual knowledge (if you graciously convince me that you liked a gift that in fact you did not, my belief about your liking will be at odds with your actual feeling) they increase the feeling of shared good will—the mutual understanding that this interaction and relationship are moving in a cooperative direction. Rituals such as greetings, gift giving, thanking, and so on establish the social grounding of an interaction.

Creating common ground is a key issue in designing spaces for the physically distant and often temporally separated participants in online conversations. In text-based interactions, the participants' primary common ground is the conversation itself. For example, people quote previous messages in their replies to create a shared context for their remarks. But, beyond automatically inserting the replied-to message, contemporary interfaces provide little support for building common ground out of the conversational record.²⁰

FIGURE 6.16

Drew Harry and Judith Donath, *Information Spaces* (2008). We designed *Information Spaces** to be a virtual space that makes social content physical. Conversations linger in the air—boxes that represent a person’s contributions to the conversation rise over his avatar’s head. In this interface, movement conveys meaning—in the scene in this picture, standing in the green area indicates agreement with a proposed idea, standing in the red, disagreement. Comments are correspondingly colored based on where the speaker was standing when making each remark. It is a virtual space composed of conversational context (Harry and Donath 2008).



Visualizing the history of a conversation—Who said what? What topics have gained much or little attention? What has been the typical rhythm of the interaction?—creates a setting that provides common ground for the ongoing conversation (see figure 6.16).

Settings are bounded; they are particular places (literally, in the physical world, and metaphorically, online). Who is privy to a particular conversation and how mores are enforced are issues of defining and defending boundaries: the repair records we saw in *History Flow* are in many cases the virtual archaeology of a skirmish over boundaries, over who defines a concept and who represents the truth. In the next chapter, we will introduce the history of contested boundaries in online discussion spaces and discuss various design approaches for maintaining strong yet also porous virtual bounds.

6 Visible Conversations: Seeing Meaning beyond Words

1. Some useful background readings from the literature on conversation and communication include Goffman 1981, especially the essay “Footing”; Bonvillian 1993; Clark 1996; Clark and Brennan 1991; Saville-Troike 2003; Tannen 2007. See Baron 2005 for comparing one type of online communication (instant messaging) with speech.
2. The anthropologist Robin Dunbar notes that nonhuman primates maintain social relationships through grooming, a process that becomes quite time-consuming as group size

grows. He has hypothesized that speech—particularly social speech and gossip—evolved as a more efficient way to manage relationships in larger and more complex groups: grooming forty or fifty comrades would leave no time for sleeping, eating, or any other activities.

3. The visualizations we discuss in this chapter make social patterns more visible and accessible; the more obvious they are, the easier it is for newcomers to understand and adapt them. But that is not always the goal: when markings of communicative competence are subtle and require one to be a long-term participant in order to be fluent, they function as a way to distinguish established users from outsiders.

4. Fernanda Viégas created all three while she was a doctoral candidate in the Sociable Media Group. *Newsgroup Crowds* was done with Marc Smith at Microsoft, *History Flow* with Martin Wattenberg at IBM, and *Themail* with members of the Sociable Media Group at MIT.

5. That said, the process of designing visualizations can itself inspire new interface ideas. Finding that you cannot depict a key pattern because the necessary data are unavailable suggests that an interface that provides those data would be useful. For example, most text archives tell us nothing about the readers—how many are there? What entries did they read most? To show these data requires an interface that logs it. We will look in greater depth at new designs for conversation interfaces in chapter 10, “Embodied Interactions.”

6. There is room here, too, for some additional and salient data. The hue of the circles, for instance, could show the ratio of a person’s initial posts to replies, which could distinguish answerers from requesters.

7. Binaries are nontext postings such as music and programs. Most of the binaries posted on Usenet are of dubious legality: pirated songs and software, and pornographic images.

8. On the function of email beyond its basic role as a conversation medium, see Ducheneaut and Bellotti 2001; Fisher et al. 2006; Mackay 1988b; Whittaker and Sidner 1996; on the rhythm of email response, see Dabbish et al. 2005; Tyler and Tang 2003; on email content, see Baron 1998; Panteli 2002; and for additional visualization approaches, see Kerr 2003; Venolia and Neustaedter 2003; Viégas and Donath 2002.

9. *Themail* used an algorithm called TF-IDF (term frequency-inverse document frequency; Salton 1988) to measure the relative importance of words. TF-IDF produces a value for the frequency of a word in a document in the context of a larger collection, which is used to determine the rarity of the word. If two words appear in the document an equal number of times, the one that appears more frequently in the contextual collection, i.e., the more common word, will have a lower value. Changing the context against which you measure rarity thus changes the value assigned each word. “*Nymphalidae*” is an uncommon word in general English usage, but common in the context of texts about butterflies. *Themail* measured the value of words in a set of messages (all the emails a person and correspondent exchanged in a month or a year) using the entire corpus of the person’s email as the context for determining rarity. This creates a personal, subjective depiction of each set of correspondences, showing how each relationship is distinct from one’s other interactions. (See also the discussion of caricature in chapter 8, “Data Portraits.”)

10. This and the following quotes are from a user study conducted by Fernanda Viégas. See Viégas 2007.

11. See Forte, Larco, and Bruckman 2009; Nov 2007; Viégas, Wattenberg, and Kushal 2004, for more on the social dynamics of Wikipedia.
12. The revision data are only a proxy for really knowing what people read. It is possible, but much more difficult, to reconstruct some notion of what has really been read through download logs (Priedhorsky et al. 2007).
13. Interruption is not necessarily rude. There are supportive interjections, overlap that is part of turn-taking, requests for clarification, etc. (Schegloff 2000; Tannen 1994).
14. Some of the physical qualities that make a leader can be reproduced online, especially in the realm of graphical interfaces. Height, for example, conveys authority, and people to whom tall avatars are given in virtual environments are treated more deferentially and act more aggressively (Yee, Bailenson, and Ducheneaut 2009). And, of course, avatars may be black or white, male or female, ugly or attractive, reproducing online the all of the status-creating distinctions of face-to-face encounters (see chapter 10, “Embodied Interactions”).
15. See, e.g., Pang and Lee 2008 on detecting sentiment and the topic of discussion; Blei and Lafferty 2006 on recognizing the topic of discussion; and Pennebaker, Mehl, and Niederhoffer 2003 on assessing personality.
16. “Celebrities” in this context are not (usually) real-world celebrities but rather celebrities in the online forum—highly active on the site, well known by all participants, etc. See Golder and Donath 2004 for an in-depth treatment of celebrities, newbies, lurkers, trolls, and other social roles in online communities. “Celebrities” and “ranters” come from this study; others might choose different names for these particular types, but the overall behavior pattern is recognizable. “Newbies” are newcomers to the site, possibly to the Internet in general. They are often unfamiliar with the accepted ways of behaving and may ask naive questions or repeat queries that have long ago been tabled by the group. “Lurkers” are people who read but seldom if ever write. Unlike the listeners in a face-to-face conversation, lurkers are often completely invisible, for many discussion technologies have no way of detecting or showing who is reading. “Trolls” are deliberate troublemakers. The prototypical troll behavior is to pretend to be a legitimate participant, but one who makes increasingly provocative comments (such as asking if using a stun gun is a good approach to feline training in a cat lovers’ group) intended to derail the conversation into rebutting these comments and arguing about whether they were ignorant but innocent, or deliberately malicious.
17. See Welser et al. 2007 for more on computational analysis and visualizing social roles in online conversation.
18. To depict this additional data, we could color the *Newsgroup Crowds* circles shown (figures 6.2 to 6.5) using a ramp of tones from red (maximum responses) to blue (no responses); moving through hues, rather than brightness (dark to light) keeps the depiction neutral. Adding a new statistic to a visualization should make salient patterns stand out better. If it instead becomes more confusing, one needs to rethink both the data choices and the design.
19. Deception is a big part of communication (Donath forthcoming). There are big lies, of course, such as saying you were at work when you were really out partying, but also small, everyday lies, like apologizing for being late because you had such trouble parking, when

really you had just gotten a late start. Polite society requires many deceptions, such as saying you liked a gift when in truth you did not; community could not exist if everyone were always relentlessly honest. Deception decreases common ground for it puts the deceived person's understanding at odds with the deceiver's knowledge.

20. See Zinman and Donath 2009 for an example of incorporating conversation history into a messaging interface to help resolve ambiguity and increase expressiveness.