

**Social Mirrors as Social Signals: Transforming Audio
into Graphics**

Karrie G. Karahalios and Tony Bergstrom

Vol. 29, No. 5
Sept./Oct. 2009

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

IEEE  computer society

© 2009 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

For more information, please see www.ieee.org/web/publications/rights/index.html.

Social Mirrors as Social Signals: Transforming Audio into Graphics

Karrie G. Karahalios and Tony Bergstrom ■ University of Illinois at Urbana-Champaign

Social visualizations are visualizations about people, for people.¹ A social mirror is a type of social visualization with three particular qualities. First, it's a third-person visualization in which people can see information about themselves within information about other people. In a sense, we can perceive "us and them." Every participant sees the same visualization. Second, subtle perceivable changes of any participant appear in the visualization as they occur in near real time. This can include capturing and visualizing subtle social behaviors such as laughing and coughing. Because of this near real-time feature, people can quickly alter their behavior—and hence the visualization—so that others will perceive them as they'd like to be perceived in this public mirror. Third, social mirrors allow for exploring group patterns and behavior through real-time experimentation, replay, annotation, and reconfiguration.

A social mirror is a specific type of visualization for group interaction. Three examples of social mirrors emphasize different motivations for visualizing vocal conversation: the power of visualization to influence conversation in real time, the addition of anonymous input into group visualization, and idea formation over time.

Although social mirrors can apply to various social situations, we illustrate them here in the context of group face-to-face gatherings. Systematic research on groups began roughly in the 1930s. Much of the fundamental research in group dynamics thus far has dealt with participant status, influence, roles, norms, and cohesion in groups. Additional areas of study include networks of friendship, group emotion, cultural and structural routines, group jargon, and group rituals.² More recently, remote-collaboration studies by Gary Olson and his colleagues focused on the psychological and technological requirements necessary

for successful collaboration at a distance.³ These studies employed field experiments with commercially available technology. The social mirror is a new concept and technology that can highlight the salient aspects of group dynamics. Social mirrors can also serve as added interaction channels, becoming abstract social cues and signals to augment group interaction.

Many instantiations of social mirrors exist; here, we use three types of social mirrors to demonstrate how they can serve as social signals. First, the *conversation clock* illustrates visualization of participant behavior; it depicts social roles, norms, and turn-taking patterns. Second, *conversation votes* illustrate the augmentation of participant visualization with anonymous voting. Individual participation as well as anonymous cues contributed by participants influence this visualization, which emphasizes social identity, roles, and group cohesion. Finally, *conversation clusters* illustrate the visualization of participant content threads over time. This social mirror traces content similarity, content formation over time, and participant contribution. (The "Related Work on Visualizations" sidebar discusses related research on visualizations and contrasts this research with our three social mirrors.)

Visualizing Participant Behavior

Aural conversation is typically ephemeral. In a conversation between four people, each person will likely remember a one-hour interaction differently, depending on his or her frame of reference. We believe there's an evolutionary reason for this, and we don't want to alter it. What would happen, however, if we captured and displayed aural conversation content abstractly? That is, what if we didn't explicitly tell the audience what was happening by capturing every word and gesture?

Related Work on Visualizations

Here, we discuss related research on visualizations and contrast this work with our three social mirrors: the conversation clock, conversation votes, and conversation clusters.

Joan DiMicco and her colleagues created a shared-group display in 2004 for small groups.¹ They used bar graphs to show participants' aggregate aural participation throughout a meeting. They labeled participants as participating, overparticipating, or underparticipating. The display influenced participants to balance their participation levels. In contrast to this work, we were interested in visualizing conversation without labeling participants' roles. We wanted to see how different participants interpreted the display and why. Our conversation clock further shows subtle nuances within group dynamics such as interruption, status roles, and mimicry, as well as how below-average speakers differed from above-average speakers in modifying behavior patterns using a social mirror.

Khaled Bachour and his colleagues recently designed a tabletop visualization with the goal of balancing participation in a task-based discussion.² For example, they asked participants to rank their favorite movies. They embedded the tabletop display with light-emitting diodes (LEDs). A specific color represented each participant. The tabletop displayed each participant's contribution as a function of the number of LEDs directly in front of his or her seating area. The aggregate representation helped balance conversation, as in the research of DiMicco and her colleagues.¹

With our conversation votes visualization, we were inspired by the research of Lee Sproull and Sara Kiesler, who had studied meeting behavior of groups in face-to-face and email scenarios. They found that in the former, higher-ranking participants (employers, managers, and so forth) spoke more than lower-ranking participants. However, when the discussion continued in a mediated channel such as email, participation of members of different ranks increased.³ The mediated channel maintained accountability of participants but decreased the risk of face-to-face humiliation or evaluation apprehension, which often promotes silence. (As Abraham Lincoln once observed: "It is

better to remain silent and be thought a fool than to speak and remove all doubt.")

Sumit Basu and his colleagues created a scalable summary of spoken conversation in a linear text interface.⁴ Viewers could zoom in for more details (for example, to see the entire speech-to-text transcript) or zoom out for a summary created via topic segmentation. In recall studies in which participants were asked to answer questions related to the transcripts, Basu and his colleagues found that the scalable summaries improved recall time compared to the full transcripts. Speech-to-text recognition is difficult. The best system today produces a 3 percent error rate with speaker training, although 20 to 30 percent error rates are more common.⁵ Performing topic segmentation on the transcript can introduce further error. Our conversation clusters approach differs from the work of Basu and his colleagues in three main ways: we don't display a complete transcript of the conversation, we use Explicit Semantic Analysis (ESA) with Wikipedia to classify keywords, and we combine the skills of humans and machines to modify our learning algorithm.

References

1. J.M. DiMicco, A. Pandolfo, and W. Bender, "Influencing Group Participation with a Shared Display," *Proc. ACM Conf. Computer Supported Cooperative Work (CSCW 04)*, ACM Press, 2004, pp. 614–623.
2. K. Bachour, F. Kaplan, and P. Dillenbourg, "Reflect: An Interactive Table for Regulating Face-to-Face Collaborative Learning," *Proc. 3rd European Conf. Technology Enhanced Learning, Times of Convergence: Technologies across Learning Contexts*, LNCS 5192, Springer, 2008, pp. 39–48.
3. L. Sproull and S. Kiesler, *Connections: New Ways of Working in the Networked Organization*, MIT Press, 1991.
4. S. Basu et al., "Scalable Summaries of Spoken Conversations," *Proc. 13th Int'l Conf. Intelligent User Interfaces (IUI 08)*, ACM Press, 2008, pp. 267–275.
5. T. Bergstrom and K. Karahalios, "Conversation Clusters: Grouping Conversation Topics through Human-Computer Dialog," *Proc. 27th Int'l Conf. Human Factors in Computing Systems (CHI 09)*, ACM Press, 2009, pp. 2349–2352.

We could transform this information into graphical language and let the audience infer a meaning. We're interested in seeing how interaction is affected when another actor is present at the table in the form of an abstract visualization, which is a graphical mirror of discourse. The conversation clock is our first attempt at realizing this goal.

In designing this social mirror, we were influenced by a previous art piece called visiphone,⁴ which was on exhibition at Siggraph in 1999. This installa-

tion represented two remote spaces as a series of abstract circles over time. Visiphone is basically a graphical display to a speakerphone. Through this string of circles, visiphone portrays conversational dominance, turn-taking, agreement, and disagreement in a conversation between two remote spaces. A common quote about this piece at Siggraph was, "It shows you what you know but don't realize that you know." (Several therapists have requested this interface for use in couples therapy.)




Figure 1. The conversation clock. The visualization renders each participant in a different color. Each rectangle’s length is proportional to the volume captured at that point in time. Dots along each ring’s circumference imply that the table is active but no one is speaking. Each ring represents one minute of conversation.

Table 1. F test and statistical-significance results for above-average and below-average speakers using the conversation clock. Statistically significant features are highlighted in bold.

Participation	Metric	F value	Significance
Above average	Time spent leading a conversation (sec./min.)	F(3, 10) = 1.29	p = 0.30
	No. of turns taken (turns/min.)	F(3, 8) = 1.19	p = 0.40
	Turn length (sec.)	F(3, 5) = 9.22	p = 0.02
Below average	Time spent leading a conversation (sec./min.)	F(3, 7) = 0.52	p = 0.70
	No. of turns taken (turns/min.)	F(3, 9) = 3.89	p = 0.05
	Turn length (sec.)	F(3, 12) = 3.32	p = 0.06

The Conversation Clock

This social mirror visualizes aural participation of up to four people around a circular table.⁵ Each participant wears a lapel microphone. The system transforms the audio captured by the microphones into an abstract form, and an overhead projector renders it onto the table. Figure 1 shows a snapshot of a conversation clock rendering.

A different color represents each person’s spoken contribution in the visualization. When someone speaks, a series of rectangular bars along the table-top’s periphery represent that person’s audio. The length of the rectangular bars is proportional to the participant’s volume: the longer the rectangular bars, the louder the audio. If no one is speak-

ing, dots appear along the ring’s circumference to indicate that the table is active and capturing audio from the microphones. The most recent conversation appears on the outermost ring. Each ring represents one minute of time. As each graphical ring completes, it animates toward the table’s center, and a new ring begins at the outermost edge.

At first glance, the conversation clock gives a historical overview of turn-taking and participation within the group. By the time the meeting is over, the conversation clock has composed a picture of the full conversation. Sociologists studying group interaction have expressed considerable interest in these visualizations. They typically use the visualizations as a log and don’t show them to participants. What happens, however, when participants see this visualization evolving in front of them in near real time?

Study Results

In our initial studies, the conversation clock visualization encouraged balanced participation between the four participants. These results were similar to those that Joan DiMicco and her colleagues found.⁶ The conversation clock extended those findings by showing that visualizing a persistent timeline as opposed to an aggregate participation history encourages balanced conversation. The studies further showed how participants alter their behavior, depending on their original contribution levels. Specifically, when dividing participants into two categories, above-average and below-average speakers, we found that although above-average speakers took approximately the same number of speaking turns with and without the conversation clock visualization, their turns taken using the conversation clock were noticeably shorter in length. Below-average speakers took more turns using the conversation clock with no noticeable difference in turn length (see Table 1 and Figure 2). Users were particularly interested in how they appeared within the group landscape. This was the case with the three social mirrors discussed in this article.

Qualitative statements about the conversation clock in post-interviews included the following:

- “It was easy to judge who [was] driving [the] conversation.”
- “I was trying to look at the circle to see whether we were balanced.”
- “I realized that I could monitor my speech patterns by watching the colors. It was interesting to train myself not to say ‘umm’ as much or pause.”
- “I noticed when you’re the one talking, you want to stop. But if you’re mid-topic, you couldn’t

stop, because you had to finish your topic. But as soon as you finished your topic, you'd shut up."

- "It became all red; should green or yellow speak next?" (asked by the author who was assigned the color yellow).

We specifically didn't label any of the graphics in the conversation clock—of the participants or their behavior in the visualization. We wanted viewers to make their own assumptions and interpretations about the visualization. In fact, when shown snapshots of visualizations, participants and nonparticipant viewers built stories around them. For example, when shown Figure 3, both nonparticipant viewers and those who had used the table inferred that the blue participant who dominated the conversation was an advisor in a meeting. This, in fact, turned out to be the case.

Similar to the findings of Fernanda Viégas and her colleagues,⁷ our study showed that participants used the visualizations as artifacts for storytelling. We had repeated requests for physical snapshots of conversations to archive specific moments such as conversations with a loved one or engagements. Beyond this, people inferred behaviors about participants.

For example, both participants and nonparticipants inferred that the pattern in Figure 4a was an interruption. Figure 4b shows what participants expected turn-taking to look like, whereas Figure 4c shows how natural turn-taking actually appears. Turn-taking is almost never clean-cut. For example, people typically signal that they're ending their turn by decreasing their volume, thereby giving others the opportunity to slowly enter the discourse.

Participants used the cues from the visualization to infer whether the conversation was heated, rational, and so forth. They used cues such as those in Figures 4d and 4e to infer who was a follower and who was a leader in the conversation. The small rectangular bursts of color were typically utterances such as "yeah," "ok," and "uh-huh." In Figure 4d, the typical assumption was that the red participant was the leader. The green participant agreed with the leader and thus was the follower. This example shows how roles and status become defined through increased use of the conversation clock. These findings are congruent with those in the literature stating, "At a verbal level, people with higher status speak more often and more loudly, are more likely to criticize, command, and interrupt others, and are spoken to more often."²

We also found indications of rapport building or affiliation via the visualizations. In typical face-




Figure 2. Results of the conversation clock study for above-average and below-average speakers. The graph shows the number of seconds per minute that above- and below-average speakers led the conversation without the visualization, during the two visualization trials, and without the visualization at the final trial.

to-face situations, this happens subconsciously. Someone perceives another person's behavior such as a gesture and then makes the same gesture. In the visualization, this was evident through mimicked volume strength and phrasing or jargon. For example, the length of the rectangles became similar for participants building rapport.

We ran the conversation clock studies in an ABCA format. That is, we conducted conversations without the visualization, then with the visualization for 10 minutes, with the visualization for 20 minutes, and finally without the visualization. When we removed the visualization, behavior slowly evolved toward the previsualization behavior. Thus, the visualization influenced behavior as long as it was visible. On the basis of previous research, we expect that if groups don't see the visualizations in real time and see them only after the conversation has ceased, the balancing behavior will disappear.





Figure 3. A snapshot of the conversation clock depicting conversational dominance. The blue participant is dominating this conversation.

Figure 4. Common conversation clock patterns: (a) concurrent speaking; (b) expected turn-taking; (c) traditional turn-taking; (d) and (e) agreement. Graphical elements such as these became building blocks for telling stories using the visualization. These snapshots were extracted from actual conversation clock trials.



We're not arguing that balanced conversation is ideal. We're arguing that social mirrors such as the conversation clock have an intrinsic power to influence group behavior. Therefore, social mirrors should be carefully designed for specific situations. This can take several iterations, as evidenced by the second visualization discussed in this article.

Augmenting Participant Visualization with Anonymous Voting

The conversation clock influenced future behavior by visualizing passive aural contributions to conversation in a manner that maintained history. This provided a common group memory, which became an additional cue in the conversation. As a next step, we explored incorporating active cues to

shape the social mirror. In this second visualization, we combined the face-to-face channel with an anonymous mediated back channel.

Conversation Votes Table

Figure 5 shows participants sitting around the conversation votes table, an extension of the conversation clock with a slightly varied visual language. Participants again sat around a table wearing lapel microphones. A different color represented each participant's audio contribution. When a participant spoke, a series of rectangular bars in that person's designated color appeared on the table. However, the length of the rectangular bars was not affected by the volume of the speaker's voice. Instead, anonymous voting by the other participants at the table determined this length. If no voting occurred, all the rectangular bars remained the same length.

Figure 6 illustrates the conversation votes setup. The central, horizontal axis depicts the current conversation. As the stream of rectangles reaches the end of the table, it shrinks and moves to the table's archival portion to make room for new dialogue. In the figure, A and B indicate the table's archival portions; they're mirror images of each other.

Various Design Iterations

The conversation votes table went through several iterations. In the first, each participant had two buttons that he or she could press anonymously—either with a hand or by using a knee against the table's underside. Pressing the green button signaled agreement with the current speaker (see Figure 7); that speaker's rectangular bars consequently lengthened, and their color increased in saturation. Pressing the red button signaled disagreement; the speaker's rectangular bars consequently shortened, and their color decreased in saturation.

Ultimately, this design was flawed. With a group of four people, using the red button to signal disagreement created discord, and the speaker receiving disagreeing votes often continued speaking until the disagreeable votes ceased or until there was agreement. In retrospect, this discord wasn't entirely surprising. With a small group, the speaker might likely become suspicious about which of the other participants are disagreeing. After our first series of studies and interviews with groups of four participants, we decided to include only one button as a signal in this social mirror. This button would signal a significant contribution in the conversation.

This original design had other ambiguities. For example, an equal number of positive and negative votes resulted in a rectangular bar that was

consistent with no votes.

Subsequent iterations of conversation votes resulted in several more design changes. Figure 8 is a snapshot of the current conversation votes visualization. We studied this version with four participants around a table, each participant having one button. We told them that pressing the button indicated a significant contribution to the conversation. When someone pressed the button, the current speaker's rectangular bars lengthened, and a dot appeared next to the narrow sides of these rectangular bars to indicate a vote had been entered. This version also allowed for concurrent conversation (such as an interruption), which was quite common. Overlapping rectangular bars depicted a concurrent conversation.

Studies of this design found that, as with the conversation clock, in unstructured discussions, participants tended toward balanced participation.⁸ Participants found the experience of this interface satisfying. This is in agreement with existing group-study literature suggesting that there appears to be more group harmony when responses are public but more accuracy when responses are private.²

This second conversation votes study used a different set of participants. Whereas participants in the first study disliked the “disagree” button, a few participants in this study suggested that we add a negative signal. Although the first study indicated that a negative vote creates discord in small groups, current ongoing research suggests that large groups could benefit from this added signal. Participants further wanted more resolution for voting and more voting parameters (agree, disagree, elaborate, I don't understand, and so on).

We performed the conversation votes studies in ABA format. The discourse was a structured debate on one of three selected topics. We asked the participants to argue their opinions to the group.⁸ Although we anticipated that below-average speakers would cast more votes, this wasn't the case. Above-average speakers cast more votes than below-average speakers. Both groups, however, reported high satisfaction levels with the table experience: they were “satisfied that all opinions were represented.” The act of seeing your vote in the interface created a record and a feeling of inclusion, even when the desired outcome wasn't achieved.

Although the conversations weren't typical brainstorming sessions, participants' experiences and decision outcomes were similar to those discovered in brainstorming studies. Measured in terms of idea quality and productivity, the results of brainstorming appear negligible. Brainstorming could lead to the impression that partici-




Figure 5. Participants sitting around the conversation votes table. The voting buttons are on the underside of the table.




Figure 6. The first iteration of the conversation votes table. A different color represents each participant around the table. When someone speaks, a series of rectangular bars in that person's color appear on the central, horizontal axis. If the other participants agree with this person's comments, their votes increase the bar's length. If they disagree, the bar shortens. As the conversation progresses, earlier portions of it are transferred to equivalent archives on each side of the horizontal axis.

pants are more productive than they really are.² Yet, participants found brainstorming satisfying. Brainstorming serves another purpose besides task productivity. It provides common ground and builds social trust as participants become acquainted. It is a low-stress activity requiring little preparation. Typically, everyone present can interject ideas—the goal being to collect several potential ideas at the initial meeting. However, it can evolve into a social situation in which no one expects a concrete result.

Figure 7. The green button from the first iteration of the conversation votes table. Pushing this button indicated agreement with the current speaker; pushing the red button indicated disagreement.



Brainstorming could fail for three main reasons. First, not everyone can speak at the same time. This is also called *blocking*. Some participants might dominate the conversation, while others remain quiet and later forget or lose their train of thought. Another reason is the *free-rider problem*. Free-riding contributions are not identifiable or even necessary. *Evaluation apprehension*, or fear of negative evaluation of remarks, can also deter effective brainstorming.

Given the problems with brainstorming, perhaps conversation votes can somewhat address conversation blocking and evaluation apprehension via additional anonymous signals. With conversation votes, participants found value in interjecting votes. They could discuss their ideas vocally, and thereby be accountable for their actions, or they could vote anonymously to mitigate evaluation apprehension.

Like the conversation clock, the conversation votes visualization becomes a social cue because it visualizes vocal behavior augmented by a voting back channel. In fact, it becomes a recursive

system because voting input further augments the visualization. Both the conversation clock and the conversation votes visualization use the audio levels captured by personal microphones to create the base visualization. They don't, however, incorporate content or context, which exist in the participants' memories (and in playback, if conversations are archived). The next social mirror, conversation clusters, intelligently incorporates spoken conversation topics into the visualization.

Visualizing Participant Content Threads over Time

Imagine a group of people gathering to discuss a design problem. Five minutes into the meeting, Charlie suggests using big red circles in the design. Lucy dismisses this remark as silly, and the meeting continues. Five minutes before the meeting ends, Lucy screams, "I've got it: big red circles!" Although this scenario is contrived, it highlights the inspirational seed at the beginning of the meeting that's repeated at the end. Can we implement a system that understands and visualizes topics over time?

The conversation clusters social mirror visualizes topics during a conversation.⁹ Participants sit around a table wearing lapel microphones; as they speak, selected keywords from their conversation appear and animate slowly on the table surface. If a specific word is repeated often, it becomes brighter and more opaque. As keywords accumulate, similar keywords form topic areas and animate to form clusters surrounded by colored membranes.

Figure 9 gives a snapshot of the conversation clusters visualization. Clusters evolve and change over time. As a cluster grows, it can separate into two clusters. Similarly, two clusters can merge.

The computer initially clusters keywords into topics. However, because the computer model of social conversation doesn't understand context, as humans do, participants can modify the model in near real time. If participants disagree with a selected computational clustering, they can move a word into or out of a cluster with their fingers. They can also create a new cluster or remove an existing one. They can do these things during or after a conversation. For example, if a participant thinks a word has been placed in the wrong cluster, he or she can touch that word and drag it out of that cluster or into a different cluster. If a participant does not think a word belongs on the table, he or she can place an outstretched hand on the table and delete that word. To add a word to the table, participants can speak that word.

The machine-learning algorithm incorporates the human corrections as they occur. This process




Figure 8. The current rendering of conversation votes. This version supports one voting button for a significant contribution, and it depicts dots at the narrow edges of the rectangular bars to indicate votes.

creates a human-computer dialogue. The hand manipulation is made possible via a touch-sensitive Mitsubishi DiamondTouch table (www.merl.com/projects/DiamondTouch).

Designing the Conversation Clusters Table

The conversation clusters table aims to

- visualize content creation within groups over time, and
- create a better topic classifier for human conversation using human-computer dialogue.

It has two visualization modes: cluster (see Figure 9) and history (see Figure 10). The history mode visualizes the progression of topics over time as the conversation progresses. The horizontal axis represents time, from left to right (while looking at the words in the proper orientation). Different colors represent different topic clusters. This view also depicts cluster splits and joins over time, as well as concurrent topics. Participants can seamlessly switch between the two modes.

To create the conversation clusters visualization, we used Nuance’s Dragon NaturallySpeaking speech recognition system (www.nuance.com/naturalspeaking) to create a transcript. Rather than focusing on creating a perfect word-for-word transcript, we highlighted keywords that we could detect with high reliability. We did this to maintain reliable accuracy as well as a certain level of privacy. Those privy to the original conversation could piece together its course from the keywords; other viewers would have more difficulty.

To create the word clusters from the speech transcript, we used Explicit Semantic Analysis (ESA) on a windowed transcript.¹⁰ ESA uses Wikipedia articles as a base collection of topics. The algorithm identifies probable topics by comparing article word distributions to our query words from the transcript. The ESA server then returns a list of articles with a weighted list of the words supporting each article. The algorithm then culls the returned word sets to avoid redundancy. Keywords form clusters if the distance metric between them falls below a threshold. In a pilot study, the Wikipedia word distributions eliminated erroneous words in our speech-to-text transcript. (More information about the algorithm is available elsewhere.⁹)

However, the current conversation clusters table has limitations. It gathers and organizes keywords and topics, but it can’t understand information that would be obvious to a human listener, such as cause and effect. At the basic level, it serves as a recall and organizational tool.




Figure 9. A cluster view of a conversation clusters table. The enclosed words are within a common similarity distance. Words outside the enclosed areas are above the similarity threshold of that topic cluster. The words appear in near real time as people speak around the table.




Figure 10. A history view of a conversation clusters table. Time progresses from left to right. Participants around the table can zoom in and out of this view to see more words from the interaction. They can also see who contributed most to a keyword. Colors of clusters change as clusters disappear, merge, and split.

User Studies

The conversation clusters social mirror is a work in progress; thus far, we’ve used it only in a pilot study. We’re designing a set of user studies to investigate these topics:

- *Content recall.* Is it easier to retrieve content from the history view or from a text search of a conversation transcript? Is the history flow an adequate representation of a meeting? Can participants recall the meeting’s general flow by looking at the history flow?
- *Cluster intelligence.* Do conversation clusters match those coded by humans?

- *Idea formation.* How do ideas form? Do individual members propose them, or does the collective group create them?
- *Visualization usage.* How do participants use the history and cluster views?
- *Ideal usage.* Is the visualization distracting during the conversation? Should the task of taking minutes at the meeting be assigned to a remote participant via conversation clusters, as opposed to the entire group? Should the visualization be on a laptop instead of on a table?
- *Effect of visualization on conversation.* Does visualization hinder conversation? Should conversation clusters meetings be structured such that the visualization isn't visible until the last 5 minutes of the gathering, at which point participants come together to cluster the discussed topics?

Our pilot study shows that the human-computer topic clustering with conversation clusters produces better results than the machine-learning algorithm alone. The human-machine dialogue occurs at two levels: first, when people add content to Wikipedia (because the algorithm uses Wikipedia to calculate topic distances), and second, when users manipulate the keywords and clusters on the conversation clusters table. The visualization helps us better understand group structure by adding another dimension with which to analyze it: content. This is relevant for several reasons. For example, existing research shows that repeated expression of an attitude can make that attitude more extreme and polarize a group.² A social mirror such as conversation clusters could highlight this phenomenon and provide insight into social influence, group cohesion, group conflict, and the road to agreement.

Further Exploration of Social Mirrors

To further explore the social mirror as a social signal, we address social mirrors through the lens of traditional group structure study. We also describe some application areas for social mirrors and explain how social mirrors evolve to become social signals and cues in our networked environment.

Social Mirrors and Traditional Group Structure

Three areas commonly critiqued in group dynamics are consequence, context and content, and group and role formation. Here, we extend these traditional themes to include visualization.

Consequence. We've seen the consequence of visualizations with the conversation clock and conversation votes. In both cases, the visualization

influenced participants to balance their conversation contributions. Those who originally spoke less overall, took more turns; those who originally spoke more, took shorter-length turns. Moreover, in the various iterations of conversation votes, we found that negative voting provoked hostility in small groups. We also found that participants wanted more parameters to express themselves in the voting process. Furthermore, our research shows that social mirrors have the potential to alter behavior through visual feedback. One serious implication of these findings is that care must be taken in the design of social mirrors. In particular, for group interaction, the social interface shouldn't "punish" participants (as it did with negative voting in one iteration).

Context and content. Conversation clusters provided the most explicit context and content of the three visualizations, with its depiction of conversation topics. Conversation votes added context, through active participant voting, that wasn't in the aural conversation. The descriptions of status, roles, and situations in the conversation clock (the example where the advisor was speaking in Figure 3) suggest that context was implied in these interactive digital artifacts. Moreover, participants altered their behavior in real time as a form of self-maintenance or self-monitoring to influence context, impression formation, and involvement level. The visualization graphics further hint at friendship or conversation networks between frequent conversation partners. This is evident in the conversation clock snapshot in Figure 4b.

Group and role formation. The social mirrors in this article don't describe group formation—that is, how the group came to be. Our future work with the conversation clock and conversation clusters will look at formation in the long term. We'll use an altered version of the conversation clock in social-skills groups with teenagers who have Asperger's syndrome. The goal is to encourage specific social behaviors such as turn-taking, varied turn lengths, and increased volume. This new visualization will define a behavior that the participants should mimic. A point (rewards) system will make the table more like a game than a chore. Suggesting group structure or altering existing meeting structures for other meeting styles could provide insights into group structure and decision formation.

So far, most of our studies have comprised at most four meetings. That isn't long enough to understand group formation, cohesion, or dis-

integration. Group structure changes over time. Although a group might begin a project with a brainstorming session, the meetings' tasks change as the project progresses. We plan to follow these dynamics with conversation clusters over longer periods of time in longitudinal studies at a local consulting firm.

Role formation is already evident in the conversation clock visualization. Participants and outside observers can see roles forming over time in the visualization. How this changes in the long term remains to be seen.

Application Areas

We've discussed these visualizations primarily in the context of meetings. But social mirrors have other application areas. They can reflect community status (whether the community is healthy and thriving with many people communicating, is dominated by one person, or is slowly losing its members and dying), the affective intensity of the group, and so on. Three other specific areas we're exploring with social mirrors are therapy sessions, group memory or archiving, and remote connections.

Therapy. In addition to our work with visiphone,⁴ we're using nuanced visualizations of voice to encourage vocalization in populations of children with speech-language delays. The participants' level of involvement in building and refining self-presentation suggests that social mirrors could be useful for additional forms of therapy.

Group memory or archiving. We don't advocate archiving everything. There are advantages to forgetting past events or remembering them in a particular fashion. However, today's technology does support constant recording and cheaper storage. If all this data is stored, how will it be retrieved?

In our studies with the conversation clock, conversation votes, and conversation clusters, salient markers or indices in the visualization often led participants to ask whether they could go back to specific parts of a conversation. With the conversation clock, participants often had moments of disagreement (marked by overlapping colors suggesting affective intensity) or moments of laughter (marked by long rectangles). With conversation votes, segments marked by votes provided entry points for exploring the conversation. In conversation clusters, people could use keywords in the history view to play back the original conversation.

The social mirrors create entry points for finding past epochs in a conversation. Participants further




Figure 11. A graph of glances at the table for the conversation clock without the visualization, during the two visualization trials, and without the visualization at the final trial. Glances to the tables increased when the visualizations were present.

requested that we provide tools for additional text and voice annotation in the visualizations. Over time, such archiving could create a transactional group memory.

Remote participation. Thus far, we've studied these visualizations in collocated settings. We're also exploring Skype visualizations (www.skype.com) of group behavior, and we've received several requests to create online versions of these visualizations for teleconferencing or inclusion of remote participants.

Limitations

Although social mirrors influence personal and group behavior, provide levels of satisfaction to users, and are a form of self-maintenance, we don't yet know how they affect group productivity. As we use these systems over the long term, we'll be able to uncover more structures in decision-making, group-think, affective influence, group cohesion, topic cohesion, and social behavior.

One concern with social mirrors that we can observe is distraction. As people sit around a table, they typically look at one another. With colorful graphics moving on a tabletop, will people focus instead on the table and keep their heads down? Will they lose the focus of the conversation while thinking about how they would like to be perceived? Although this question is still open, we have some indications that social mirrors don't require 100 percent of users' attention.

We videotaped each social-mirrors study and coded it for glances to the table during interaction. Figure 11 shows a graph of glances for the conversation clock. When the visualizations were present, glances to the table increased. However, this depended on the participants' role. If person A was speaking to person B, neither A nor B looked at the table. Persons C and D, however, made quick glances at the tabletop. In interviews and surveys, participants didn't report that the table display disturbed them. But further study is necessary.

We're also restructuring the notion of the meeting. For example, with the conversation clusters table, we propose showing the visualization five minutes before the meeting's end. In this way, meeting participants could end the meeting by grouping


the keywords into clusters after discussing them, to establish common ground. Another suggestion is to show the social mirror only to remote participants. Remote participants often feel marginalized in meetings where they're the minority. This would give them the ability to contribute actively.

Signals and Cues

Marc Hauser has distinguished between communication *signals* and *cues*.¹¹ (Social-signaling theory uses various designations for these concepts, but we use Marc Hauser's definitions in this article.) Cues are unintentional or given unconsciously. Examples include skin color, sneezes, and accents. A signal is an intentional transmission of information. For example, someone might try to lead or guide a meeting by speaking frequently and loudly to imply status, or someone might say he or she played golf with Bill Gates. Given these basic definitions, deception is possible. Although faking height in face-to-face interaction might be difficult, people can, with study and practice, fake a British accent.

The social mirrors discussed in this article form a hybrid of signals and cues. Many of the features rendered are cues. In traditional meetings without visualizations, the speaker often doesn't notice he or she is taking few turns or speaking only at the beginning of the meeting. The conversation clock, however, would highlight this behavior as a cue. The registering of votes in conversation votes and the removal of conversation topics in conversation clusters, however, are overt signals.

We created these visualizations in concert with the face-to-face channel. Thus, the social mirrors become another channel for interaction (or a back channel) and, in the process, become a signal that influences interaction.

Writing transformed thought formation, as people transitioned from an oral to a literate culture. Visualizations further alter existing oral culture into graphical language. This transformation will have further implications for group, networked, and social behavior. Social mirrors mark the beginning of this transformation. 

References

1. K. Karahalios and F. Viégas, "Social Visualization: Exploring Text, Audio, and Video Interactions," *Proc. 24th Int'l Conf. Human Factors in Computing Systems (CHI 06)*, ACM Press, 2006, pp. 1667-1670.
2. J.M. Levine and R.L. Moreland, eds., *Small Groups: Key Readings*, Psychology Press, 2006.

3. G.M. Olson, J.S. Olson, and T. Pease, *Basic Group Behavior*, Now Publishers, to be published in 2009.
4. J. Donath, K. Karahalios, and F.B. Viégas, "Visiphone: Connecting Domestic Spaces with Audio," *Proc. Int'l Conf. Auditory Display (ICAD 00)*, Int'l Community for Auditory Display, 2000, www.icad.org/websiteV2.0/Conferences/ICAD2000/PDFs/Visiphone.pdf.
5. T. Bergstrom and K. Karahalios, "Seeing More: Visualizing Audio Cues," *Proc. 11th IFIP Int'l Conf. Human-Computer Interaction (Interact 07)*, Int'l Federation for Information Processing, 2007, pp. 29-42.
6. J.M. DiMicco, A. Pandolfo, and W. Bender, "Influencing Group Participation with a Shared Display," *Proc. ACM Conf. Computer Supported Cooperative Work (CSCW 04)*, ACM Press, 2004, pp. 614-623.
7. F. Viégas et al., "Digital Artifacts for Remembering and Storytelling: Posthistory and Social Network Fragments," *Proc. 37th Hawaii Int'l Conf. System Sciences (HICSS 04)*, IEEE CS Press, 2004, p. 40109.1.
8. T. Bergstrom and K. Karahalios, "Vote and Be Heard: Adding Back-Channel Cues to Social Mirrors," to be published in *Proc. 12th IFIP Int'l Conf. Human-Computer Interaction (Interact 09)*, Int'l Federation for Information Processing, 2009.
9. T. Bergstrom and K. Karahalios, "Conversation Clusters: Grouping Conversation Topics through Human-Computer Dialog," *Proc. 27th Int'l Conf. Human Factors in Computing Systems (CHI 09)*, ACM Press, 2009, pp. 2349-2352.
10. E. Gabrilovich and S. Markovitch, "Computing Semantic Relatedness Using Wikipedia-Based Explicit Semantic Analysis," *Proc. 20th Int'l Joint Conf. Artificial Intelligence (IJCAI 07)*, IJCAI, 2007, pp. 1606-1611; www.ijcai.org/papers07/Papers/IJCAI07-259.pdf.
11. M. Hauser. *The Evolution of Communication*, MIT Press, 1996.

Karrie G. Karahalios is an assistant professor in the Department of Computer Science at the University of Illinois at Urbana-Champaign, where she runs the Social Spaces Group. Her research focuses on the interaction between people and the social cues they perceive in networked electronic spaces. Karahalios has a PhD in media arts and science from the Massachusetts Institute of Technology. She's a member of the IEEE and ACM. Contact her at kkarahal@cs.uiuc.edu.

Tony Bergstrom is a PhD candidate in the Department of Computer Science at the University of Illinois at Urbana-Champaign, where he's a member of the Social Spaces Group. His research focuses on interfaces influencing conversation through visualization. He has an MS in computer science from the University of Illinois at Urbana-Champaign. Contact him at abergst2@cs.uiuc.edu.