Visualization of Audio: A social tool for face-to-face groups

Joan Morris DiMicco
Sun Microsystems Inc.
1 Network Drive
Burlington, MA 01803
joan.dimicco@sun.com

Katherine J. Hollenbach
MIT Media Lab
20 Ames St
Cambridge, MA 02139
kjhollen@mit.edu

CHI 2006, April 22–27, 2006, Montréal, Québec, Canada.
ACM 1-59593-298-4/06/0004.

Abstract
This position paper presents a visualization system for viewing and reviewing turn-taking patterns in a face-to-face meeting, used as a way to gain insight into the social dynamics of a group.

Introduction

Social psychologists suggest that one way a group can improve its interaction and consequently its productivity is by having a high-level understanding of its emotional and social interaction [5]. Furthermore, a group can gain insight into its interaction by viewing an aggregate report of the its behavior [1]. Therefore, we postulate that an automated method for collecting and visualizing social dynamic information can assist a group in understanding and improving its interaction.

To explore this hypothesis, we developed the application Second Messenger, a platform for providing feedback to groups during their meetings about participation levels and after meetings about more general patterns in turn-taking and overlapped speech. This paper presents these visualizations.

The Second Messenger Visualizations

To capture the turn-taking dynamics during a meeting, each person using Second Messenger wears a noise-canceling microphone that detects when he or she is speaking. Second Messenger aggregates these time-stamped moments of individual speech to create a record of who spoke when in a meeting and uses this log to create visualizations, to be viewed either real-time during the meeting, or afterwards as a replay. The sound levels on the microphones are collected in parallel, enabling Second Messenger to detect simultaneous speech.

The application has five visualizations that display the same data in different visual forms. Figures 1 through 7 show these visualizations, illustrating the same moment in a conversation of six people.

Our hypothesis is that when a group sees this histogram, the most talkative will allow the quietest to have the floor and contribute to the discussion.

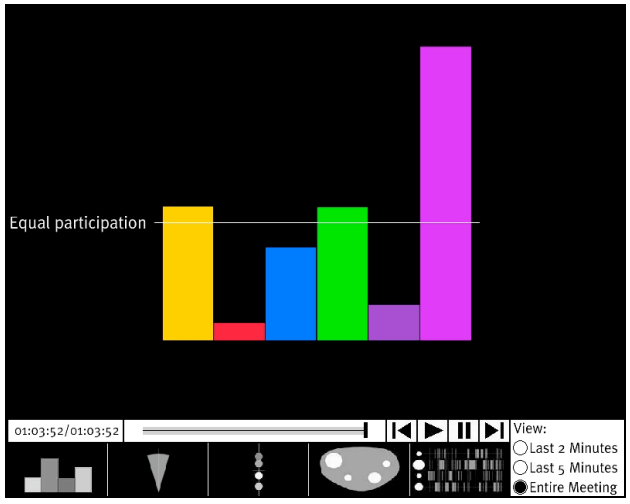

This screenshot includes the application window, with playback controls, visualization selection menu, and temporal window selector.

figure 1. Histogram visualization

The first visualization is the Histogram (figure 1). The vertical bars on the screen represent how much each person has spoken in relation to others and adjust dynamically while a meeting occurs. This visualization was used in our previous work that examined how real-time feedback influences participation levels [3]. In that experiment, we found that showing a group this histogram led those at the highest levels of participation to decrease the amount they spoke. While equal participation in meetings is not our goal, we view this finding, that a simple display can change behavior, as encouragement for continuing to pursue our goal of improving group interaction.

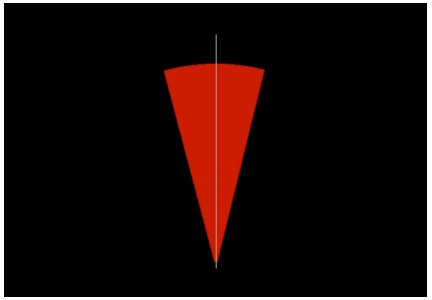

As the spread in group participation increases or decreases over time, the width of the fan grows or shrinks accordingly.

figure 2. Fan visualization

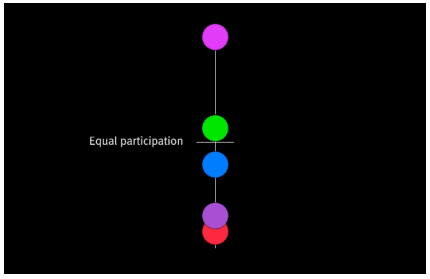

When the floor control switches from one person to another, one circle will move down while another will move to the top.

figure 3. Bouncing Balls

Clicking on each of the balls toggles their identifying color to monochrome. When monochrome, the further the ball is from the horizon, the more red it appears.

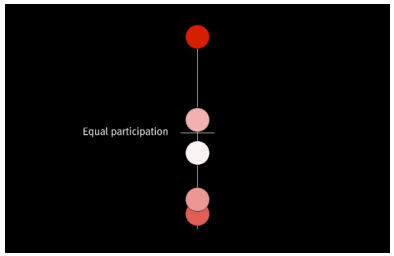

figure 4. Bouncing Balls, anonymous

Using a pointer, each of the circles can be arranged to mimic the physical arrangement of the face-to-face group.

The grey shape around the circles is meant to reflect and emphasize the intimate nature of a face-to-face group, in contrast with the other visualizations which emphasize comparison and evaluation.

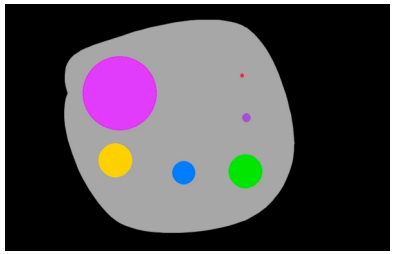

figure 5. Group Circle visualization

By clicking on each of the circles, the proportion of overlap speech is highlighted.

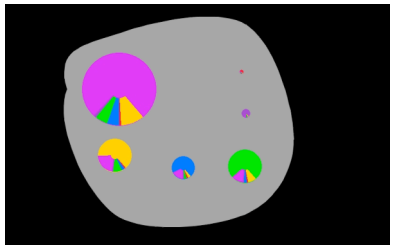

figure 6. Group Circle, with overlap

The timeline can show the entire meeting (as shown) or a two-minute or five-minute window, providing a detailed view of the meeting's turn-switching.

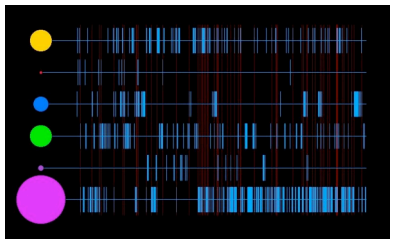

figure 7. Timeline visualization

In the Fan visualization (figure 2), a single fan image represents the spread in participation between members, calculated as the standard deviation between participation values. This image removes any individual identity and treats the group as one unit, where an extreme imbalance is signaled by an abnormally wide fan. In contrast with the other screens, this removes the possibility of singling out one individual as an outlier in participation and potentially encourages the group to work together to be inclusive.

The next visualization, the Bouncing Ball (figure 3), represents each person as a ball of a fixed diameter where the ball's vertical position reflects the individual participation level. When viewed with a limited time window (a two or five minute range), this view provides a way of seeing who currently has the floor (the upward moving circle) and who is falling away from the center of the conversation (the downward moving circles). Figure 4 shows the Bouncing Ball in "anonymous mode" where the circles are colored in monochrome, becoming more a saturated red as they move to the extremes. As with the fan, this is an attempt at visualizing the imbalance without pointing toward single individuals.

The Group Circle (figure 5) represents each group member as a circle that grows and shrinks in proportion to that person's participation level. Figure 6 shows the Group Circle where each circle is highlighting speaker overlap. Within each person's circle, pie slices reflect how much each *other* person spoke *while* that person was speaking. Overlapped speech, while commonly thought of as interrupting, also includes instances of simultaneous laughter, short verbal affirmations, and side comments. Noticeable differences in overlap levels

within a group can indicate a unique relationship one or more of the individuals has with the group. For example, when one person in the group is an outsider, this person will most likely have a proportionately smaller amount of overlap on his/her circle, reflecting the more tentative interaction a group has with a new or visiting member.

The final visualization is the Timeline (figure 7) that reveals who is speaking at each moment in the meeting. The circles down the left side of the screen represent the individuals and the horizontal lines extending from the circles have vertical blue bars at the moments when that individual spoke. To highlight moments of overlapping speech, transparent vertical red lines are drawn where more than one person spoke. Using this view in replay mode, a user can watch a meeting's timeline scroll past and see who spoke exactly when and who gained the floor at moments of overlapping speech.

Related Work

The field of CSCW includes numerous examples of systems designed for summarizing meetings. Two examples of automated collection of audio information include: Kristjansson et al. [6] who developed a method for extracting meeting keywords through voice recognition to annotate a video of a meeting; and Chen [2] who built a system to provide feedback on who is speaking, gesturing, and fidgeting in a remote classroom to help an instructor feel the "pulse of the classroom." As with our work, they chose simple visualizations of activity to make different behaviors more salient. In addition to CSCW, our work on Second Messenger draws from the field of social visualization, for techniques on conveying social meaning [7].

Implications & Conclusion

The visualizations of audio data in Second Messenger provide a platform for testing our hypothesis that providing groups this information will benefit their interaction. We have used Second Messenger in several behavioral experiments and given the tool to a number of real-world groups [4]. The experiments found that a display revealing real-time participation levels causes those at the highest levels of participation to *decrease* the amount they speak. Reviewing the timeline visualization after a meeting causes those who spoke the least to *increase* the amount they speak in a subsequent discussion; real-time visualizations did not produce this change. Additionally, after groups review their turn-taking patterns, the experiments found that they alter their information-sharing strategies in later discussions.

These findings move towards demonstrating how additional social information provided to groups, in the form of abstract visualizations, changes behavior, perceptions of the interactions, and the process of group decision-making.

Acknowledgements

This work was funded by the MIT Media Lab's research consortia.

References

[1] Bales, R.F. Interaction process analysis: a method for the study of small groups. Addison-Wesley, Reading, MA, 1950.

[2] Chen, M. Visualizing the Pulse of a Classroom, In Proceedings of the ACM International Conference on Multimedia (MM'03) (Berkeley, CA). 2003.

[3] DiMicco, J.M., Pandolfo, A. and Bender, W. Influencing Group Participation with a Shared Display, In Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW04) (Chicago, IL). 2004.

[4] DiMicco, J.M. Changing Small Group Interaction through Visual Reflections of Social Behavior. PhD Thesis, MIT Media Lab, Massachusetts Institute of Technology, Cambridge, MA, 2005.

[5] Hackman, J.R. Groups That Work (And Those That Don't): Creating conditions for effective teamwork. Jossey-Bass, San Francisco, 1990.

[6] Kristjansson, T., Ramesh, P., Huang, T.S. and Juang, B.H. A Unified Structure-Based Framework for Indexing and Gisting of Meetings, In Proceedings of the IEEE International Conference on Multimedia Computing and Systems (ICMCS) (Florence, Italy). 1999, 572-577.

[7] Viégas, F. and Donath, J. Chat Circles, In Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 99) (Pittsburgh, PA). 1999.