Overt or Subtle? Supporting Group Conversations with Automatically Targeted Directives

Gianluca Schiavo[1&2] **Alessandro Cappelletti**[1] **Eleonora Mencarini**[2&3]
Oliviero Stock[1] **Massimo Zancanaro**[1]

[1] FBK-irst
Via Sommarive, 18
38123 Trento, Italy

[2] Trento RISE
Via Sommarive, 18
38123 Trento, Italy

[3] University of Trento
Via Belenzani, 12
38122 Trento, Italy

{gschiavo | cappelle | mencarini | stock | zancana}@fbk.eu

ABSTRACT

In this paper, we present a system that acts as an automatic facilitator by supporting the flow of communication in a group conversation activity. The system monitors the group members' non-verbal behavior and promotes balanced participation, giving targeted directives to the participants through peripheral displays. We describe an initial study to compare two ways of influencing participants' social dynamics: overt directives, explicit recommendations of social actions displayed in the form of text; or subtle directives, where the same recommendations are provided in an implicit manner. Our study indicates that, when the participants understand how the implicit messages work, the subtle facilitation is regarded as more useful than the overt one and it is considered to more positively influence the group behavior.

Author Keywords

Conversation support; Visual Attention; Social dynamics; Implicit interaction; Persuasive Technologies.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Although face-to-face group discussions are very common in most workplaces as well as in informal settings, effective interaction is not always easy [7]. There has been a large amount of research into multimodal support of group conversations and many different attempts have been made to influence group activities providing real-time feedback on the group's social dynamics. We specifically aim to contribute to this research by considering a different perspective: the use of a directive approach based on subtle and minimally obtrusive directives to change the group behavior.

IUI'14, February 24–27, 2014, Haifa, Israel.

ACM 978-1-4503-2184-6/14/02...$15.00.
http://dx.doi.org/10.1145/2557500.2557507

We use the term "directive" with reference to the illocutionary force [22] of a communicative act to have the explicit intention to induce the receiver to perform a certain action.

In this work, we present an ambient intelligence system that provides a minimal form of facilitation to foster a group of people to balance their individual participation during a face-to-face conversation. The system targets specific participants on the basis of the group behavior and gives directives to them. Specifically, we aim at providing evidence that subtle directives may be not less effective but definitely more acceptable and less obtrusive than overt directive messages.

Our approach is inspired by one of the key functions of a facilitator, namely to "*ensure that all team members contribute*" [5]. The importance of involving all the members does not always mean that all participants need to participate equally. Here, we take the simplified view of fostering equal participation for the sake of supporting group dynamics where all the points of view are equally considered. Indeed, in some contexts, such as brainstorming, the equal participations of all the members has been shown to improve the final outcome both in terms of quantity and originality of the ideas generated [17].

We adopt a "calm technology" [28] stance: designing interfaces that remain in the background providing information in a calm and unobtrusive manner. In our case we want the technology never to become prominent and to distract from the intellectual and social activity. Ultimately, we want to investigate a specific example of a persuasive technology [9] where the behavior change promoted by the technology is specifically achieved through social influence, but not through coercion, deception and with a minimal obtrusiveness.

We designed a system that monitors the behavior of the participants and assesses the level of participation of each individual. The system purposefully intervenes in the dynamics of the group by displaying the directives on personal displays embedded in a piece of furniture. The tablet devices act as peripheral displays [1] in the sense that they are not intended as the primary focus of the participants' activity, but they allow participants to be

aware of information without being overburdened or distracted from their main activity.

In designing the visualizations which provide the subtle directives, we took inspiration from the information decoration approach, which aims at creating a balance between aesthetic and informational quality, and where classical aesthetical factors such as ambiguity and repetition are used to achieve interesting images that convey information in pleasant manner [8].

In a study conducted in a real-world setting, we compared the subtle directives, with and without the explanation about the information portrayed, and a control condition consisting of overt directives conveying the same communicative meaning in the form of text. The results suggest that, when the participants understand how the implicit messages work, the subtle facilitation is regarded as more useful than the overt one and it is considered to have a more substantial positive influence for changing the group behavior.

Although preliminary, these results suggest that the approach of subtle directives provided in an aesthetical and unobtrusive manner might be an effective way for persuading a group to change their behavior.

RELATED WORK

Within the domain of technology that aims at supporting group activities, many different attempts have been made to influence group conversations by providing real-time feedback on the social dynamics. Kim et al. [10] have used sociometric sensors to monitor communication patterns and other social signals during team activity, reporting a graphical representation of group dynamics to the members themselves. For example, Meeting Mediator [10] is a system that detects social interactions using sociometric badges and provides visual feedback on a mobile screen with the aim of enhancing group collaboration in tasks where balanced participation is desirable. According to the authors, cooperative behaviors can be promoted by visualizing social signals to the group members in order to increase their awareness of the communication patterns.

Situated displays and tabletops have been augmented with sensors to increase and support group activities taking place in the device proximity. For example, DiMicco et al. [6] have investigated peripheral displays that visualize the amount of participation of the member of a small group conversation. Specifically, the authors investigated the effect in group meetings of a shared display that visualizes real-time information of how much each person has spoken in relation to the others. Participation was considered in terms of vocal activity measured using close-talk microphones. This information was used to visualize each member's participation with the purpose of stimulating individual reflections on the on-going activity and harnessing social collaboration. The system was evaluated both in structured tasks (i.e. information-sharing and decision tasks) and in real-world situations (e.g. work-meetings) showing that such technology can influence groups' behavior towards a more balanced level of participation. Specifically, the system effectively made the over-participants to decrease the amount they spoke, but the behavior of under-participants did not changed.

Other researchers have investigated the use of tabletop technologies to influence social dynamics during group conversations [26, 29, 13, 14, 4]. These interactive tables are combined with sensors for monitoring individual and group behavior, such as vocal activity, head orientation and body fidgeting. The multimodal information is used to generate unobtrusive visual feedback about the dynamics of the group conversation with the aim of supporting with various degrees of awareness, the social interactions of the users involved in the task. The visualizations include information about each individual's behavior, including quantitative measures of the physical interaction with the system [14], analysis of the vocal activity, such as the personal cumulative speaking time since the beginning of the meeting or the duration of the current turn [4], and also patterns of visual attention, such as the attention given to other participants as a listener or the attention received while speaking [26]. The visualizations are generally presented in a graphical representation displayed in real-time on the table's surface, using projected display [29, 26] or a LED board embedded on the surface [4]. The visualizations provide the participants with feedback on their performance [26, 4] or with material for supporting and sustaining the activity [29], and their applications have been explored in domains such as cultural heritage [29] and collaborative learning [12, 13].

Research in the area of smart meeting systems provides another set of examples of technologies that act to support participants in meetings by automatically analyzing the group behavior [15, 20]. For example, Pianesi et al. [18] have shown that multimodal information from meeting participants can be used to model the group dynamics and to generate post-activity summary reports that can help participants to improve their awareness on the meeting performance. Other studies have investigated how technology can facilitate turn-taking and pace of a conversation in meetings. For example, Time Aura [11] is a desktop application designed to help people to control their pace, while giving an oral presentation. Time Aura shows an overview of the presentation's structure, real-time information and suggestions about the progress of the task and feedback on speaker's performance. Different feedback modalities have been investigated in the case of Occhialini et al. [16], who presented an interactive lighting system for time management during meetings. The system was composed of halogen lights used to display a dynamic pattern on the walls, giving peripheral information to the speaker about the meeting progress. Variables of the lights, such as direction, color and intensity, were manipulated to reflect the meeting progress while blinks notified the

speaker of the approaching end of the available presentation time.

Besides using graphical and explicit feedback, more subtle and indirect modalities have been investigated. Rogers et al. [21] explored the use of ambient displays to nudge behavioral change in people (i.e. whether to take the stairs or the elevator) using subtle and abstract modalities of feedback. In a further study, Balaam and colleagues [2] showed how a multi-user public display can enhance interactional synchrony by displaying peripheral subtle feedback about users' nonverbal behavior related to rapport (including simultaneous movements, posture matching, eye contact and related back channel responses). A Wizard-of-Oz study indicated that the participants showed significantly more interactional synchrony and more coordination when their nonverbal manifestations of rapport were amplified by the ambient display. According to the authors, the findings suggest that social dynamics, like rapport, can be leveraged by the technology to support group behavior, without requiring a direct and exclusive interaction with the users. Indeed, these studies showed that the effect of the technology was not always explicitly perceived by the users: the displays indirectly influenced participants' nonverbal behavior, while at the same time participants were not aware of this change.

Taken together, these studies show that dynamic feedback, presented both explicitly (e.g. through text, graphs and real-time statistics) or implicitly (e.g. using peripheral animations), can potentially foster members experience and affect group behavior.

In the present study, we present a prototype aimed at facilitating balanced conversation in groups of four people. Our system is similar to previous work in the way it adopts sensors to monitor social dynamics. The main difference is in the manner in which the system addresses the participants and influences group activity. In this regard, instead of providing real-time feedback, which requires self-reflection, we wanted to investigate a directive style of intervention: the system proactively sends messages with an illocutionary purpose to the group participants.

During the deployment of such technology, we compared two ways of influencing participants' social dynamics, namely using an overt and a more subtle style of intervention.

THE TECHNOLOGY

The system has the form factor of a small piece of furniture, such as a small table. The present prototype includes 4 Kinect cameras and 4 10-inch tablet displays which are placed in the center of the group organized in a circle (Figure 1). Each participant sits on a colored chair placed in pre-determined positions at a distance of about 1.2 m from the table. The chair's color is used to identify the different participants in the communication strategies (discussed below). The positions of the chairs were determined to make sure the perceptual capabilities of the Kinect are optimized, while allowing for a good perspective and suitable distance from the display. At present, each display is meant to be visible only to the person in front of it.

In the next two sections, we will describe in details the mechanisms for monitoring the social dynamics in the group and the visual interventions (subtle vs. overt).

Scene Analysis

The visual scene is analyzed using the 4 Kinect cameras. Each camera monitors a single participant by detecting the position of the head (as provided by a module in Microsoft Kinect API) and tracking where they are looking, using an algorithm that considers the relative positions of the other participants in the setting. Four situations are modeled: whether the person's head is directed toward the participant on the left, on the right, directly opposite, or toward no one in particular. This latter case denotes the state where the person is not looking at other participants (specifically, when the person is looking at the table). The yaw angle, used to discriminate between front and lateral orientations (left or right) and the tile angle were 35° and 5°, respectively (Figure 2). These values were based on both empirical tests and previous studies in the literature [26]. In this work, we assume that the head orientation of a person reasonably approximates the gaze and the focus of attention, as has been shown in previous studies in meeting settings with 4 people [24].

Figure 1. A group during a conversation (left) and a zoom on the prototype (right).

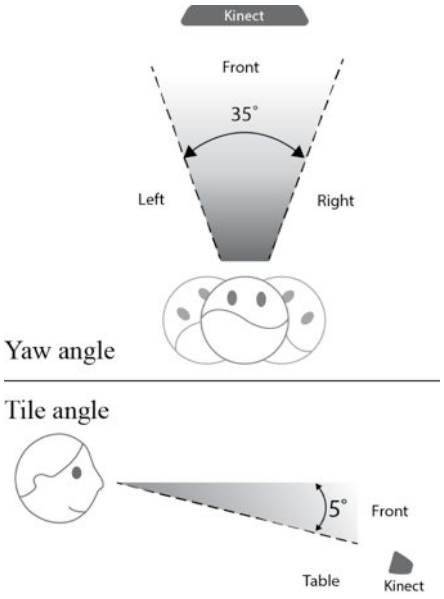

Figure 2. Diagram of the angles used to discriminate between head directions.

From the information of who is looking at whom, considered as an evidence of a participant's focus of attention, the system creates the following aggregate measures:

Visual Attraction (VAT) is defined for each participant as the relative amount of visual attention received by the other participants in a given time interval;

The *Visual Distribution Index* (VDI) in a given interval of time is defined as the degree to which each member is evenly looked at. The index is based on the minimum value of Visual Attraction in the same interval (the number 4 in the equation is given by the fact that our system at present expects exactly 4 participants)

$$\text{VDI} = \frac{\min(VAT)}{1/4}$$

VDI values range between 1 and 0, where 1 defines a balanced attention distribution, where all participants are equally looked at, and 0 denotes an unbalanced distribution, where at least one participant is not looked at by anyone else. In the current prototype, VDI is updated every 20 seconds. VAT and VDI were then used by the system to trigger the visual intervention.

Visual intervention

As explained above, we aimed to experiment with directive interventions: that is, our system is supposed to suggest specific behaviors to the participants by delivering directive messages.

In order to compare overt, language-based, directives with subtle, aesthetically pleasant ones, we focused on a single strategy aimed at fostering inclusion of the group members with low level of participation.

The subtle directive messages are presented by means of visualizations designed using the principle of information decoration [8]. That is, the informational content of the message is presented in a form that highlights aesthetical qualities in such a way that it can be "calmly" presented in the environment. Ambiguity and repetitions are classical means of achieving interesting representations. In our case, our messages consisted of an animation of bubbles, whose colors varied to address each different participant, while the movement of the bubbles was meant to "suggest" the performativity (in this case, directing the attention of the group toward the person who is under-participating in the conversation).

By default, the personal screen of each participant displays an animation of several bubbles of the same colors of their own chair moving toward them. This animation is not a directive per se but it represents an idle state of the system.

When the VDI drops below a minimum threshold (for the current prototype, the value has been set to 0.2, based on previous pilot studies with 4 groups), the participant with the lowest VAT in the last 20-second interval is selected as the target of the group who should be more involved in the conversation. The communicative strategy to achieve this goal is implemented as follows (Figure 3):

- The target participant sees on his/her display the bubble of his/her own color moving circularly and becoming more transparent;
- The other participants see on their displays all the bubbles changing to the color of the target chair and moving toward the target position.

As a control condition for the evaluation study, we also implemented an overt communication strategy as follow (Figure 4):

- The selected participant sees on his/her display one of the following textual messages (randomly selected to reduce repetition): "Come on! You too can contribute", "You're not participating enough! Say something!" and "What do you think about it?"
- Each other participant sees on his/her display one of the following messages (again, randomly selected to reduce repetition): "You should include everyone in the conversation!", "Maybe the person on the [red/ green/yellow/purple] chair wants to say something!", "What does the person on the [red/green/yellow/purple] chair think?".

The text appeared on the screen using fading animation, remaining for 20 seconds before disappearing. The default screen for the control condition is a neutral background with no textual messages.

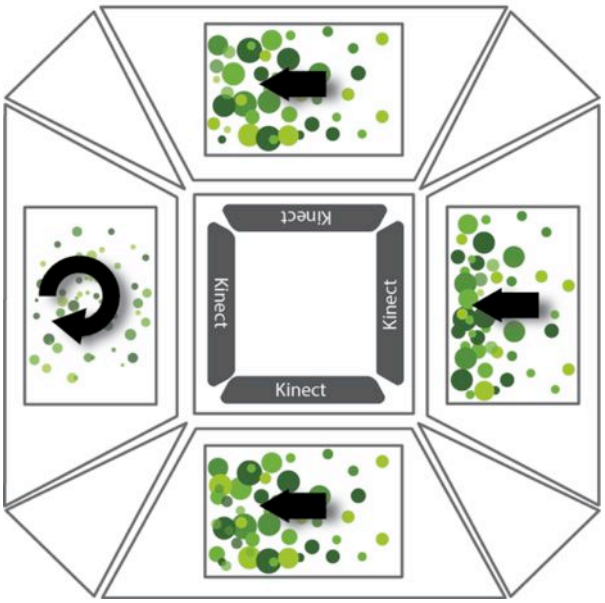

Figure 3. Visualization of the subtle strategy. The participant on the left (identified by the green color) is the current target. The arrows indicate direction of the movement.

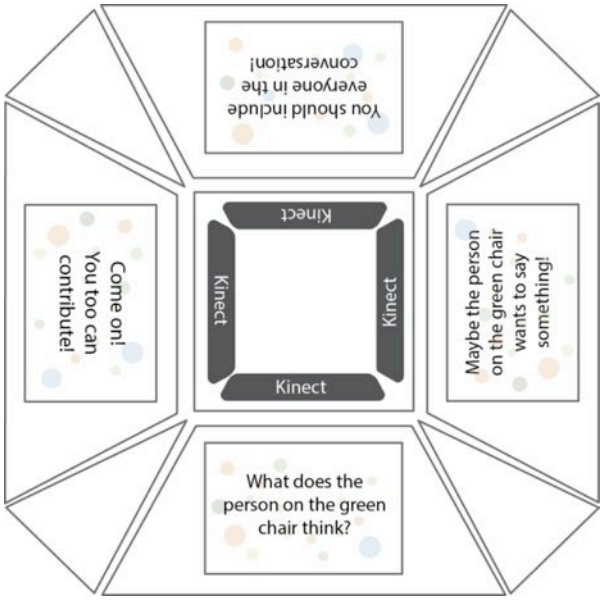

Figure 4. Visualization of the overt strategy. The participant on the left (the target) receives directive messages to participate. Other participants receive directives to include the target in the conversation.

EVALUATION STUDY

The evaluation study was organized as a quasi-experiment in the context of an ecological setting: a large science fair taking place in Trento (Italy). Visitors were invited to participate in a brainstorming session (as a way of inducing a conversation). Groups were either formed at the moment or composed of friends. The impossibility of controlling all the variables (e.g. group composition, previous exposure to the explanation, and so on) characterizes the study as a quasi-experiment [23]. Furthermore, a *mixed methods* approach [25] has been used, in which both quantitative and qualitative data are collected.

The study used a between-subject design with 3 conditions:

- *Subtle directives with prior explanation (S_W)*. In this condition the group was informed before starting the task about the system's behavior and the information portrayed.
- *Subtle directives without any explanation (S_WO)*. The system used subtle visual directives, but the meaning was not previously explained to the group.
- *Overt directives (O)*. The system showed directives in the form of text.

Hypothesis

In comparing the subtle with the overt messages, we expect that the latter would be quite effective in changing the group dynamics and eventually lead to a more balanced participation. Thus, our first hypothesis (H1) is that the subtle condition would perform as well as the overt condition, in particular when the participants understand the meaning of the visualizations (condition *S_W*).

On the other hand, we expect (our second hypothesis, H2) that the subtle visualizations (conditions *S_W* and *S_WO*) would be less invasive than the overt ones (condition *O*). In this respect, the difference between the two subtle conditions has been introduced to evaluate the effect of an explicit understanding of the meaning of the visualization respectively on the impact of the intervention on the group dynamics and on its obtrusiveness.

Task

The system was evaluated in the context of a science fair, where brainstorming on topics related to the fair is a well-known technique for encouraging informal learning. Participants were asked to brainstorm ideas on how to improve ecological practices in daily life, such as, reducing waste and being more sustainable. Some cards with stimuli for discussion were prepared on the topic to help participants in developing their arguments. In the course of the experiment itself, only one group actually made use of them.

Procedure

At the beginning of each session, a group of participant was invited to sit and was randomly assigned to one of the three conditions. The experimenter instructed the group on the brainstorming activity and on the topic of their discussion. The participants were then asked to brainstorm for 10 minutes. All the sessions were recorded with written consent of the participants.

A short assessment session followed. First, the participants were asked to fill an individual questionnaire with 10 items that took about 5 minutes to complete (Table 1). The questions were about their experience (3 items), the disturbance of the system (2 items) and the perceived cohesiveness of the group (5-item scale adapted form the Price and Muller Work Group Cohesion Index [19]). Then, a semi-structured group interview (of around 10 minutes) was conducted to understand participants' experience with the technology. Finally participants were debriefed about the study and about all the study conditions.

Results of the study

Fifteen (15) groups, a total of 60 people (38 women and 24 men), participated in the study. Five (5) groups experienced the subtle visual intervention with prior explanation, 6 groups the subtle visual intervention without explanation and 4 groups experienced the system with the overt messages.

Assessment of the visual scene analysis

The data collected in the study was also used to assess the technical capabilities of the system in terms of monitoring functionalities. Here we report some figures about the system performance.

The Kinect cameras tracked the participants for an average of 60% (subtle condition with and without prior explanation = 57% and 58% respectively. Overt condition = 69%).

To assess the reliability of our system and head orientation tracking, three one minute meeting fragments were selected (one for each condition). For these fragments, an expert coder manually annotated head orientations for each participant using a video annotation tool. Cohen's kappa was used as a measure of reliability between automatic and manual annotations. The accuracy for classification of head direction was 68% and the resulting kappa was .67, showing good agreement.

Furthermore, in order to assess the capability of the system to correctly identify the under-participating individuals (who are the target of the of a specific type of system's interventions), two independent observers were asked to perform exactly this task by observing the videos of the conversations (excluding audio). These observers watched 20 clips (each clip was 20 seconds long) and labeled the under-participating person(s). The agreement between the observers' scores and the estimations provided by the system was good, with kappa of .65. These results show that our system provides a reasonably good estimation of the group members' participation in an ecological setting.

Distribution of Attention

According to our first hypothesis (H1), we would expect a difference in the overall balancing of the distribution of attention among conditions.

In order to measure the balancing of the group attentional behavior, we estimated the equality of the distribution of attention at the group level by defining an index inspired by the Gini coefficient. The Gini coefficient [27] is a measure of equality that sums, over all the group members, the deviation of each member from an equal distribution, normalized by the maximum possible value of the deviation. Similar metrics inspired by the Gini coefficient have already been adopted for measuring participation equality in group conversations using speaking time [6, 12, 26] and visual attention [26].

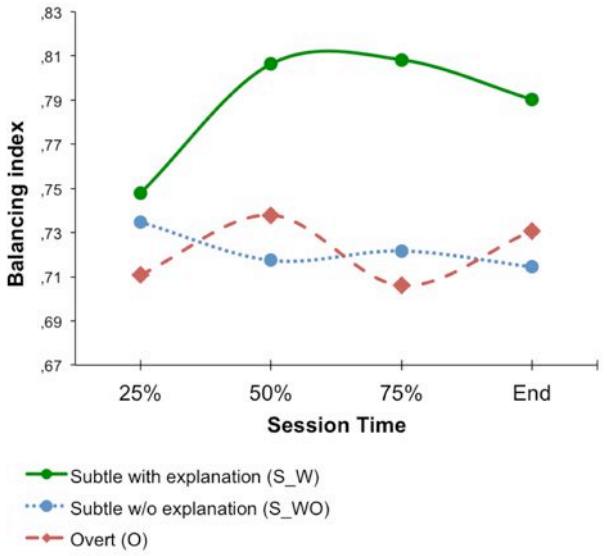

Figure 5. Average values of the Balancing index by conditions across session time.

We define the *Balancing index* as follow:

$$\text{Balancing index} = 1 - {}^{2}/_{3} \times \sum_{i} |\text{VAT}_i - 25\%|$$

The *Balancing index* represents to what extent the amount of attention received by each participant was equally distributed, considering the *Visual Attraction* (VAT) of all group participants and not only the minimum value. The index ranges between 0 (unequal distribution) and 1 (equal distribution): the more the distribution of attention deviates from an equal distribution, the lower will be the resulting value of the index.

Figure 5 shows the average distribution of the *Balancing index* across the session time. In *S_W* the index of equal attention is generally higher after the first half of the sessions (M= 0.79), while in the other conditions the index remains, relatively stable across time (subtle without explanation & overt, M= 0.72). Even though the distribution was in the expected direction, the difference was not statistically significant. These results partially support H1.

Items	Subtle with explanation ***S_W*** (*N*=20)	Subtle without explanation ***S_WO*** (*N*=24)	Overt ***O*** (*N*=16)	F and p values
I found the information displayed useful	2.90 (*0.97*)	1.17 (*0.48*)	2.13 (*1.09*)	$F_{2,57}$= 23.86 **p < .01**
I think the system had an influence on the discussion	2.90 (*1.12*)	1.83 (*0.96*)	2.50 (*1.15*)	$F_{2,57}$= 5.60 **p < .01**
I think the system had a negative / positive influence	3.65 (*0.88*)	2.88 (*0.68*)	3.13 (*0.62*)	$F_{2,57}$= 6.10 **p < .01**
I was distracted by the displays	2.20 (*1.60*)	2.13 (*1.50*)	2.31 (*1.69*)	$F_{2,57}$= 0.13 p > .05
I was distracted by the Kinect camera	1.60 (*0.94*)	1.50 (*0.83*)	1.69 (*1.01*)	$F_{2,57}$= 0.24 p > .05
Cohesion	4.08 (*0.57*)	4.09 (*0.67*)	4.07 (*0.57*)	$F_{2,57}$= 0.04 p > .05

Table 1. Responses (means and standard deviations) to the questions about the display on a 5-point Likert scale from 1 (*Not al all*) to 5 (*Very much*). *N*=60

Questionnaire results

All the groups reported high scores on the cohesion dimension (M= 4.08, SD= 0.57; the Cronbach's alpha for this scale was 0.81). Participants generally indicated that they had enough ideas on the brainstorming topic (M= 3.0, SD= 0.8) and that they were satisfied with the group work (M= 3.8, SD= 0.8). They also were not disturbed by the presence of the Kinect cameras (M= 1.58, SD= 0.91).

Regarding the perceived usefulness of the system (Table 1), an ANOVA indicates a statistically significant difference in the scores for the three conditions ($F_{2,57}$=23.86; p<.01). Post-hoc pairwise comparisons (Tukey HSD correction) showed that participants in the subtle condition with prior explanation perceived the system as more useful compared to the conditions with overt (p< .01) and subtle messages with no explanation (p< .01). This result supports H1.

In addition, participants in condition *S_W* (who had a prior explanation of the subtle intervention) reported that the system had stronger ($F_{2,57}$=5.60; p< .01) and more positive ($F_{2,57}$=6.1; p< .01) influence on the discussion compared to *S_WO* (p< .01), but no significant differences were observed compared to overt messages (p> .05). Again, this partially provides support for H1.

Participants in all conditions reported that the displays were generally slightly distracting (M= 2.20, SD= 1.11) but no significant differences across conditions were observed.

Interview results

The interviews indicate that the participants in condition *S_WO* noticed that the content changed over time, but they were not able to give meaning to the animation. Only one group correctly guessed that the bubble movements were related to their conversation, but they could not conceive an explanation of how the system worked.

The following excerpt from the interview illustrates this:

The system has surely some relation with our conversation, but we didn't know the meaning, so I would not say that it was useful. (from Group 5, Participant 2).

Mostly participants in this condition tended to ignore the display during the conversation. As an effect, they regarded the animation as meaningless and uninteresting; they related to it as a sort of screensaver.

In condition *O*, the participants clearly understood that the purpose of the system was to include everyone in the conversation. They were aware that the system showed different messages to the other participants at different times but they were not aware of what conditions led to the personalized messages. They found the messages displayed somehow distracting and intrusive, not really helpful for their activity (which is in line with the questionnaire scores). This attitude supports part of H2, as expressed in the following statement:

[The system] was useless to the communication and distracting. When I read the text I lost the thread of the conversation. (G4, P4).

In condition *S_W*, some participants (at least one in each group) reported that they actively used the display to monitor the conversation and to adjust their behavior accordingly. They looked at the displays purposely, monitoring the direction and the color of the bubbles to be aware of the conversation flow:

I used the system for making the conversation more balanced: I spoke more when I saw the bubbles moving in

circle and I tried to include the person suggested by the bubbles when they move toward someone. (G11, P3).

Yet, other participants felt uncomfortable when the system targeted them as the under-participant:

I'm a shy person. When I saw the bubbles of my color fading on the display I felt even more excluded. (G14, P4)

Most people, however, indicated that the animation adequately represented the group dynamics (e.g. targeting the person who was not really in the conversation) even though they did not constantly devote their attention to the displays. This may partially support both H1 and H2, even if it opens some issues, which will be discussed below.

DISCUSSION AND FUTURE WORK

Our research hypotheses were that (H1) the subtle directives might equally or slightly less effective than the overt ones and that would be less obtrusive (H2).

The data presented above shows some evidence, that for H1, a stronger hypothesis may hold: namely, that subtle directives may be even more effective than the overt messages provided when the participants have a prior understanding of the former's meaning. This evidence is suggested by an improvement of the Gini-like distribution of the visual attention which consistently happened in the last part of the task for the groups in the condition of subtle directives with prior explanations (yet, it is fair to note that the effect is not strong enough to be statistically significant, probably because of the small number of groups involved).

Still, the subjective measures from the questionnaires and the interviews confirmed that the explained subtle condition is perceived as more useful than the other two.

The distributions for the groups in the other two conditions do not seem to show an improvement of the balancing toward the end of the task. Indeed, the subjective measures from the questionnaires and the post-task interviews seem to indicate that this lack of improvement is for different reasons: the subtle messages (without explanation) tend to be considered meaningless and uninteresting while the language-based directives are distracting and obtrusive and not perceived as helpful for the activity.

Regarding hypothesis H2, we did not find any strong evidence from the questionnaire since all the interventions were scored as moderately distracting, but the findings from the post-task interviews support the hypothesis that the overt intervention was perceived as more distracting and not useful.

The study had some limitations: the present prototype is limited to four users and improvements in accuracy are advisable; the monitoring of the group dynamics was simple and a single communication strategy and visualization style were tested; last but not least, both the number of groups involved and the time spent in the task were relatively small. Still, the study was conducted in an ecological situation where the system was evaluated in a real-world environment.

Although still preliminary and not conclusive, our study suggests the effectiveness of subtle directives, which are non-verbal and aesthetically pleasant illocutionary acts, for changing the dynamics among a group of people.

This approach represents a different stance with respect to the extensive use of awareness feedback systems [6, 26, 4]. In our case, the system provides a directive illocutionary message that is intended to make the recipient perform a certain action. Our initial hypothesis was that a subtle directive may be not less effective than an overt directive but our results seem to exceed our hypothesis and suggest that it may work even better than the overt one when the meaning of the visualization is understood. In the study, this condition was assured by a prior explanation of the system's animation.

The idea of illocutionary force of a communicative expression is that it results from the intention of the message originator to bring about a certain state of affairs by intervening (linguistically, in the original concept) on another agent. The expression is meant to cause the receiver to perform an action. A condition is that the message originator is confident that his/her expression is recognized as a communicative message, and that the recipient has the means for understanding it (e.g. lexicon, grammar). Of course the originator may assume there can possibly be unintended ambiguity in the message (for instance the recipient may wrongly interpret the request for getting information about the time in "do you know what time is it" and answer "yes"). But the capability to correctly interpret the message by the recipient, aside from ambiguity, is assumed. At the other extreme, we have directive codes, like the illocutionary meaning of the crossroad lights which are unambiguously understood and shared.

In this work, as said, we wanted to explore subtle, minimally obtrusive influence on participants' behavior. The specific visual intervention is not the focus of attention in this paper, it is just an example of a class of modes. We noted that subtle directives might not meet the precondition of being understood as a communication mode, and, as consequence, mostly no meaning was sought for the message by the recipients. This is the reason we have introduced in the design experiment a condition where the subtle intervention was explained. This allows us to consider the case in principle, without imposing a conventional code, which tends to be very simple and without any ambiguity.

As future work, we will investigate different visualization styles to find calm, unobtrusive and aesthetically pleasant modes. We will seek among these modes those that lead to an understanding that there is a message which grant that the meaning is normally well understood, even when there

is the possibility of ambiguous interpretation. We also plan to investigate the effect of different modalities for visualization: while at present we just experiment with personal display, we will also use shared or partially shared screens.

More in general, we want to pursue the study of subtle illocutionary forces. In particular, we will better explore different types of illocutionary forces in communicative strategies in order to provide a general framework for subtle intervention.

An interesting topic that emerged from the interviews is that some participants felt uncomfortable with the system because it can socially reveal their under-participation. Although this feeling is probably not only specific to directive systems (it may be applied as well to awareness systems), it might impact on acceptability of social persuasive systems in the long term. In future work, we will focus on this aspect by investigating under what condition a participant may feel uncomfortable and how we can design solutions that alleviate this aspect.

Finally yet importantly, this kind of technology might raise the issue of being manipulative by influencing participants' behavior. In our work, we are consistent with the approach discussed by Berdichevsky and Neuenschwander [3] since the intended outcome of our system is not unethical. Yet, we recognize that the method of using subtle messages may be at risk of improper use. What we aim for is an unobtrusive approach that reduces users' cognitive demand rather than an approach in which the user is not aware of being manipulated. However, it might be difficult to draw a clear line to separate the two. We believe that further investigations on the general topic of subtle illocutionary force may help in defining a general framework for better understanding this aspect.

CONCLUSION

This paper discussed a system that acts as an automatic facilitator by supporting the flow of communication in a group conversation activity. We presented a prototype that unobtrusively tracks group participation, acts in the periphery of users' attention and can targets specific group members. Previous work has mainly focused on using technology to provide feedback on the group behavior. With this study, we specifically aimed to contribute to this research area by adopting a different perspective: the use of a directive approach to change the group behavior realized through a subtle, non-distracting and aesthetically pleasant manner.

We acknowledge the limitations of this study, including the lack of variability in the type of visualizations and in the choice of different strategies. Yet, these simplifications were necessary to set up a controlled study in a real-world setting. Our main finding is that subtle directives can influence group behavior and are more positively accepted, compared to overt, language-based directives. These results contribute to the debate on persuasive technologies and are specifically meant to lead to novel applications to support collocated groups.

ACKNOWLEDGMENTS

This work was conducted as part of the PER TE project, supported by Trento RISE. We thank Radu-Laurentiu Vieriu, Sergey Tulyakov and Nicu Sebe for helping in the design and implementation of the head tracker algorithm. We want also to thank all the participants who took part in the study.

REFERENCES

1. Bakker, S., van den Hoven, E., and Eggen, B. Acting by hand: Informing interaction design for the periphery of people's attention. *Interacting with Computers* 24, 3 (2012), 119–130.
2. Balaam, M., Fitzpatrick, G., Good, J., and Harris, E. Enhancing interactional synchrony with an ambient display. Proceedings of *CHI 2011, SIGCHI Conference on Human Factors in Computing Systems*, ACM (2011), 867–876.
3. Berdichevsky, D. and Neuenschwander, E. Toward an ethics of persuasive technology. *Communications of the ACM* 42, 5 (1999), 51–58.
4. Bergstrom, T. and Karahalios, K. Social Mirrors as Social Signals: Transforming Audio into Graphics. *IEEE Computer Graphics and Applications* 29, 5 (2009), 22–32.
5. Burns, G. The secrets of team facilitation. *Training and Development* 49 (1995), 46–52.
6. DiMicco, J.M., Hollenbach, K.J., Pandolfo, A., and Bender, W. The impact of increased awareness while face-to-face. *Human Computer Interaction* 22, 1 (2007), 47–96.
7. Doyle, M. and Straus, D. How to make meetings work. Berkley Trade (1993).
8. Eggen, B. and Mensvoort, K.V. Making Sense of What Is Going on 'Around': Designing Environmental Awareness Information Displays. In P. Markopoulos, B.D. Ruyter and W. Mackay, eds., *Awareness Systems*. Springer London, 2009, 99–124.
9. Fogg, B. J. Persuasive Technology: Using Computers to Change What We Think and Do, Morgan Kaufmann Publishers (2003).
10. Kim, T., Hinds, P., and Pentland, A. Awareness as an antidote to distance: Making distributed groups cooperative and consistent. Proceedings of *CSCW 2012, ACM conference on Computer Supported Cooperative Work*, ACM (2012), 1237–1246.
11. Mamykina, L., Mynatt, E., and Terry, M.A. Time Aura: interfaces for pacing. Proceedings of *CHI 2001, SIGCHI*

Conference on Human Factors in Computing Systems, ACM (2001), 144–151.

12. Martinez, R., Kay, J., Wallace, J.R., and Yacef, K. Modelling Symmetry of Activity as an Indicator of Collocated Group Collaboration. In J.A. Konstan, R. Conejo, J.L. Marzo and N. Oliver, eds., *User Modeling, Adaption and Personalization*. Springer Berlin Heidelberg, (2011), 207–218.
13. Martinez-Maldonado, R., Kay, J., Yacef, K., and Schwendimann, B. An Interactive Teacher's Dashboard for Monitoring Groups in a Multi-tabletop Learning Environment. In S.A. Cerri, W.J. Clancey, G. Papadourakis and K. Panourgia, eds., *Intelligent Tutoring Systems*. Springer Berlin Heidelberg, (2012), 482–492.
14. Martinez-Maldonado, R., Dimitriadis, Y., Martinez-Monés, A., Kay, J., and Yacef, K. Capturing and analyzing verbal and physical collaborative learning interactions at an enriched interactive tabletop. *International Journal of Computer-Supported Collaborative Learning* 8, 4 (2013), 455–485.
15. Nijholt, A., op den Akker, R., and Heylen, D. Meetings and meeting modeling in smart environments. *AI & Society* 20, 2 (2006), 202–220.
16. Occhialini, V., Essen, H. van, and Eggen, B. Design and Evaluation of an Ambient Display to Support Time Management during Meetings. In P. Campos, N. Graham, J. Jorge, N. Nunes, P. Palanque and M. Winckler, eds. *Human-Computer Interaction – INTERACT 2011*. Springer Berlin Heidelberg, (2011), 263–280.
17. Oxley, N.L., Dzindolet, M.T., and Paulus, P.B. The effects of facilitators on the performance of brainstorming groups. *Journal of Social Behavior & Personality* 11, 4 (1996), 633–646.
18. Pianesi, F., Zancanaro, M., Not, E., Leonardi, C., Falcon, V., and Lepri, B. Multimodal support to group dynamics. *Personal Ubiquitous Computing* 12, 3 (2008), 181–195.
19. Price, J. L., & Mueller, C. W. Handbook of organizational measurement. Pitman Publishing (1986).
20. Rienks, R., Nijholt, A., and Barthelmess, P. Pro-active meeting assistants: attention please! *AI & Society* 23, 2 (2009), 213–231.
21. Rogers, Y., Hazlewood, W.R., Marshall, P., Dalton, N., and Hertrich, S. Ambient influence: can twinkly lights lure and abstract representations trigger behavioral change? Proceedings of *Ubicomp 2010, 12th ACM international conference on Ubiquitous computing*, ACM (2010), 261–270.
22. Searle, J.R. A classification of illocutionary acts. *Language in society* 5, 1 (1976), 1–23.
23. Shadish, W., Cook, T., and Campbell, D. Experimental and Quasi-Experimental Designs for Generalized Causal Inference, Houghton-Mifflin (2002).
24. Stiefelhagen, R., Yang, J., and Waibel, A. Modeling focus of attention for meeting indexing based on multiple cues. *IEEE Transactions on Neural Networks* 13, 4 (2002), 928–938.
25. Teddlie, C.B. and Tashakkori, A. Foundations of Mixed Methods Research: Integrating Quantitative and Qualitative Approaches in the Social and Behavioral Sciences. Sage Publications, (2009).
26. Terken, J. and Sturm, J. Multimodal support for social dynamics in co-located meetings. *Personal Ubiquitous Computing* 14, 8 (2010), 703–714.
27. Weisband, S.P., Schneider, S.K., and Connolly, T. Computer-mediated communication and social information: Status salience and status differences. *Academy of Management Journal* 38, 4 (1995), 1124–1151.
28. Weiser, M. and Brown, J.S. Designing calm technology. PowerGrid Journal 1, 1 (1996), 75–85.
29. Zancanaro, M., Stock, O., Tomasini, D., and Pianesi, F. A socially aware persuasive system for supporting conversations at the museum café. Proceedings of *IUI 2011*, *International Conference on Intelligent User Interfaces*, ACM (2011), 395–398.