

Persuasion at the Museum Café: Initial Evaluation of a Tabletop Display Influencing Group Conversation

Cesare Rocchi, Oliviero Stock, Massimo Zancanaro, Fabio Pianesi and Daniel Tomasini¹

Abstract. A café table is a traditional setting for conversation. Tabletop displays may have an active role in this connection. In particular for a museum scenario, conversation after the visit is important for a joint elaboration of a small group visit experience. We propose the museum café as the location to introduce a tabletop display meant to foster and support conversation about the visit. The goal of the system is to influence the development of the conversation by adopting persuasion techniques. We describe a system that monitors the conversation among the visitors and dynamically shows visual stimuli on the table surface. An initial formative evaluation is conducted through a series of qualitative user studies.

1 INTRODUCTION

In museum scenarios, informal conversations among small group of visitors play a fundamental role in the learning process, as ethnographic studies have clearly demonstrated [4]. We take this as the inspiration for this work. How can technology induce people to entertain a conversation about their experience at the museum and help sustain it? This question is closely related to the work of Fogg [2], which highlighted the potential of computers as persuasive tools that can influence people's behaviour, also in an educational entertainment scenario like the one we propose.

Most of the current technology for museum visits addresses the single user [3]; people, however, tend to visit a cultural site with families, groups of friends, etc. Petrelli and Not report that 45% of the visitors go in organized groups [4]. Mobile guides and kiosks thus are in risks of hampering rather than fostering conversation. We propose a novel aspect: technological tools that provide support after the visit, when visitors can have a conversation about their experience. In particular we investigate a tabletop application placed in the museum café specifically designed to influence the subject of the conversation and the behaviour of the group.

The table is instrumented with sensors and its top surface is used as medium to display persuasive messages aimed at influencing the conversation of the group. Conversation is tracked through word spotting [5], and the group's behaviour is monitored. Reasoning about the overall conversation configuration and the visit to the museum permits to drive the system actions: the

system chooses specific presentation strategies that lead to specific output on the tabletop.

The scenario, at the museum café after the visit, includes three phases:

- a) a phase where the system promotes a conversation about the museum visit experience with the goal of shifting the group discussion into a specific topic of the cultural experience;
- b) a phase that supports conversation by providing content appropriate to the specific topic being discussed and the state of the conversation;
- c) a phase where one member or the whole group explicitly seek further information about some cultural heritage topics by interacting with the system.

In the present paper we concentrate on the first two phases, where there is no explicit input to the system by the participants; the system observation of the non mediated participants' interaction is used as a sort of implicit input. When people entertain a conversation on a topic not related to the visit, the system tries to influence the conversation by attracting them towards visit-related topics, with techniques reminiscent of the tradition of advertising. When the conversation is about a museum topic, it supports it by proposing relevant material, also drawing on information about their visit history.

In the following, after a short review of related works, we report about an initial Wizard of Oz user study investigating subjects' reactions to the tabletop display during a conversation after a museum visit. The study explores the effect of a number of communicative strategies exploited by the system, which are borrowed from semiotics and advertisement techniques. The analysis of video recorded data and post-study interviews has helped defining the technological requirements. The actual architecture implementing those requirements is then described in the following section. Presentation strategies are specifically in focus in the following part, with the accent put on the persuasive connotation. We then describe the present implementation and briefly discuss it.

2 RELATED WORKS

The system proposed in this paper has some affinity with a peripheral display [6] in that the table is not central to the attention of the group and people may look at it only occasionally. Yet peripheral displays are normally used as secondary sources of information, separate from a user's primary, focal task [7] and are usually meant to have a passive role and just aim at making users aware of easily graspable

¹ FBK-irst, 38100 Trento, Italy.

Email: {roccoli/stock/zancana/pianesi/tomasini}@fbk.eu

information such as weather or stock graphics. On the opposite, our system actively monitors the group's behavior with the intention to induce specific behavior by displaying information when contextually appropriate.

There have been many studies on the display of information through social tools. For example, the Tangible Bit project was about conveying information to increase awareness of people's presence and activity [8]. Another example is Groupcast, a wall projected office application that creates informal interaction opportunities by displaying mutual interest to people passing by [9]. Drift is an interactive table that displays an aerial photo of England through a hole, to foster interpretation and engagement [10]. Qualitative observations showed that people got engaged by interacting with the system and narrating about the places spotted. Hello Wall is a digital wall made of a grid of lights [11]: depending on people distance, the wall changes communicative function (ambient, notification, interaction). Abstract light patterns convey information about mood, presence and crowdedness.

DiMicco and Bender [12] have experimented with a system that, by monitoring working groups, presents information about relational behavior in the form of graphical displays on a tabletop device, to affect group behavior.

A similar approach was pursued by Sturm and colleagues [13] who used a tabletop device as a peripheral display aiming at the same self-regulatory effect as discussed above. In their approach, they display not only the speaking time but also the gaze behavior of their participants. Their results show a similar effect of Dimicco and colleagues for what concern the speaking behavior and no effect on gaze behavior.

Kim and colleagues [14] used a portable device called the sociometric badge to monitor speaking activity and other social signals in a team. They report a graphical representation of the group behavior on a private display. Their results showed a reduction of the overlapping speech but not a significant increase in solo speech.

All these approaches are based on the idea that reflection on one's own behavior may bring to rational decisions about behavior changing [15]. Usually these systems are applied in a team-work scenario where each participant is motivated to achieve his/her goal, e.g. a successful meeting and/or a well accepted personal appearance. Their approach, focused on balancing the contributions of the participants, has been proven to be effective in reducing the involvement of dominant participants but not in increasing the participation of the less active ones.

We propose a different approach: our system intends to affect the group behavior by presenting on a shared interface (namely the café table) contextually appropriate visual material in a novel way, reminiscent of the tradition of advertisements: attention catching, evocative and cognitively stimulating.

Our approach is motivated also by studies in the field of persuasive technology [2]. Fogg identified seven strategies for persuasive technology tools:

- Reduction: making something complex appear simpler.
- Tunneling: demand to an expert.
- Tailoring: providing relevant information.
- Suggestion: act at the right time with a message.
- Self-monitoring: tracking the desired behavior

- Surveillance: publicly observe one's behavior
- Conditioning: reinforce target behavior with positive "reward".

With reference to the above strategies, we propose a system that features suggestion and conditioning. Suggestion is based on interventions at the right time to maximize the effectiveness of the persuasive message. A suggestion-based technology actively induces someone to do something she might not have done otherwise. In our case, the system shows stimuli meant to support the current activity of the group (a conversation about the visit experience) or favor a behavioral change (make some participant more active). Conditioning is based on the provision of positive feedback to favor the persistence of an already occurring behavior. This strategy is usually adopted when the system aims at supporting an ongoing conversation. Our system also uses a tailoring strategy, which appropriately selects content according to the topic currently discussed. The stimuli presented by this strategy are related to the current topic discussed, in a way similar to recommender systems [16].

Data about the conversation are processed by the system to output stimuli that realize suggestion, conditioning and tailoring strategies. In a nutshell, instead of revealing to the group the social dynamics and requiring them to take into account the information and act rationally to achieve a given meeting goal, we aim at directly influencing (modifying or sustaining) the behaviour of the group.

3 THE INITIAL STUDY: WIZARD OF OZ

A Wizard of Oz experiment was initially performed to study the reaction of the users to an active table in the museum café (see [17] for the details). We hypothesized that data available to the system are: images and texts about the exhibition, profiles and visit's logs for each visitor and an automatic speech system able to understand the topic of the conversation.

In this study three groups of 4 people were invited in our lab to visit a reconstruction of the "Cycle of the Months" frescoes in Torre Aquila, (Castello del Buonconsiglio), Trento, Italy. Subjects were given a four-page booklet to help them during the visit and were told that the purpose of the study was to test the content of the booklet. After the visit people were conducted to another room and were invited to sit at a table while waiting for the experimenter to come back. The wizard, located in another room, monitored the group behaviour and controlled the presentation of visual stimuli projected onto the table. After the study, an experimenter debriefed the group about the real purpose of the study and conducted a semi-structured interview aimed at eliciting subjective impressions.

Recorded sessions and interviews have been analysed addressing the following questions:

- Did a stimulus catch the attention of one or more users?
- Did a particular system action (e.g. zooming on an image) favour the change of a topic?
- Did a graphical effect upset the users?

The questions of the interview addressed the role of images and words, the density of stimuli displayed and the conceptualization of the system's behaviour in general.

From the observations and the interviews, it resulted that the system was recognized as a useful tool to wrap up a visit, especially in case people were not acquainted with the exhibition. Subjects also reported the feeling that the table sometimes ‘follows’ the conversation and tries to propose new hints. They also said to be upset in case of weird behaviour, especially when the image supporting the conversation disappears. All the groups reported that when the discussion of a topic was languishing they used the stimuli on the table to start a new conversation. Yet graphic-intensive effects like pulsing and flashing have been considered too upsetting, especially when there is an ongoing conversation.



Figure 1. A snapshot of the Augmented Café Table

4 THE SYSTEM

Starting from the insight from the WoZ we developed the Augmented Café Table, a system that analyses the conversation of a group of people around the table and presents a set of stimuli in order to change or sustain the conversation. The Augmented Café Table is a tabletop display with the form factor of a café table (see Figure 1). At present, the interface is top-projected, for future releases we will experiment with back-projection and multi-touch capabilities. In the final scenario, people are sitting at a museum’s café table after having visited the frescoes. The current system targets the “Cycle of the Months” frescoes referred to above. The “desired” topics for conversation are the frescoes themselves, and the set of stimuli exploited by the system includes images and videos of the frescoes or of related details thereof, as well as short sentences relevant in the domain.

The system employs a set of microphones to capture the users’ conversation, which is analysed using a keyword spotter. Knowledge about the behavior of each individual in the museum is also traced using the visit’s logs from a multimedia guides. Logs provide information about the exhibits visited, and the amount of information the guide provided. As said above, the system has two roles:

a) the system promotes conversation about the museum visit experience trying to shift the discussion towards a specific visit-related topic;

b) the system provides contents appropriate to the current topic of discussion and the state of the conversation.

The interface displays visual stimuli such as floating words and pictures meant to be cues for the conversation, whereby the group can discuss ideas, share impressions, exchange opinions and, in general, get along with the spirit of the visit.

The system is organized along three modules: perception, interpretation and presentation (see Figure 3).

The perception module receives and processes data from sensors. The first type of data relates to voice activity, captured by microphones. The perception module includes a Voice Activity Detector and a keyword spotter, which recognizes words uttered during the conversation. The output of the keyword spotter is a series of words, with attached a measure of confidence indicating the reliability of the result. A second type of data comes from the visual scene and where face detection mechanism deals with visual attention toward the table. The input from four webcams is processed by a Haar-based head detector [18]. The visual attention of each participant is estimated by calculating changes in the proportion of the bounding box of the face: when the box is vertically squeezed, the system detects a visual attention toward the table.

The interpretation module analyses the state of the conversation, which is modelled along two dimensions: state of the group and state of the table. The first dimension relates to the content actually discussed and the behaviour of the participants around the table, while the second deals with the actions of the table itself. One of the features of the conversation state is the topic currently discussed. Topic detection is based on a semantic network as in Figure 2. Each node of the network is a topic, which represent relevant concepts related to the visit. Our network includes one topic for each panel of the fresco, plus general topics which are shared by more than one panel, e.g. life in the middle age or information about the restoration. Each topic has attached a set of keywords, which describe the node.

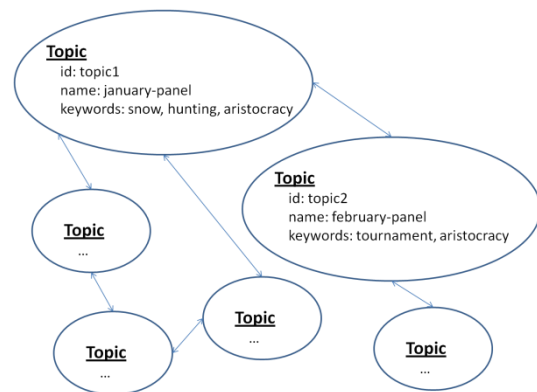


Figure 2. An excerpt of the topic network.

Such model allows the system to compute topics’ “connectedness”. This information is then processed by the strategy selector to dynamically propose new topics to be discussed. In doing so, the system takes into account the history of the conversation. The elements of the history are topics, with attached information about the duration (how long a topic has been discussed) and the list of speakers who contributed to the

discussion. In case the system has not recognised anything pertaining to one of the expected topics, it just records information about participants' activity and the topic of the conversation is marked as "out of domain" assuming that the group is conversing on something unrelated to their museum experience. The history of the conversation is used by the interpretation module by considering also its evolution in a given time window:

- level of participation . E.g., a person has been too much or too little active.
- conversation development. E.g., the recent conversation has been "jumpy", i.e. recently discussed topics are unrelated one to the other.
- topic coverage. E.g., the current topic has already been discussed extensively.

The information discussed above is cumulatively used by the system to trigger a reasoning mechanism which selects an appropriate presentation strategy to be realized on the tabletop surface (see below). Beside the current state, the choice of a presentation strategy is also conditioned by the presentation currently displayed on the surface and the history of strategies previously used.

For example, to support continuation of the conversation on the same or related topics, and enforce the cohesion of the tabletop dynamics, the system will reason on the topic network and both on the history of the conversation and the history of its own presentations.

Animation of images and words is a key characteristics of the stimuli displayed on the table. Motion captures attention and is easier to identify in the periphery than color and shape [19]. A proper timing of animations is indeed of paramount importance since the onset of motion is more effective at capturing attention than motion itself [20]. We think that some of the design dimensions commonly adopted by peripheral displays are useful to structure the presentation layer. We consider three dimensions: data representation, notification and transition.

Data representation refers to the way stimuli are shown and the potential impact they can have on the conversation. Since we want to immediately influence the development of the conversation we use images and words as possible stimuli to foster and support a conversation about the visit. The goal of this choice is twofold: to allow focusing on particular aspects of the painting and to foster the visitors' interpretative engagement. By interpretive engagement we mean the attitude of asking questions like: "What is this? What is my experience about it? What can I share with others about it?"

Notification relates to the dynamics the system adopts to show visual stimuli. The way stimuli appear and move is meant to catch the attention and potentially change the behavior of the group. A stimulus appearing or getting larger can have diverse effects on people's perception: it can be simply change blind (a person viewing the visual scene does not detect large changes in the scene) attention grabbing, it can increase awareness, or even directly demand physical action by the user. The notification layer can be thought as a sort of rhetoric of information presentation. The display actions we implemented are meant to obtain different effects in dependence of the state of the

conversation. For example, a way to change the topic is to grab the attention of a group member on a detail. In this case the notification has to be clear and indicate a passage: for example only one stimulus is visible and the image is progressively enlarged. On the other hand, if the goal is to introduce a topic related to the current one, the change can be smooth—e.g., a new detail is displayed via a slow fade among already present stimuli. Notification is related also to the visual patterns which objects can be organized into. For example a strategy to notify that certain objects are related exploits the metaphor of spatial proximity, that is more related objects are located closer. In our system we implemented a notification pattern that aligns stimuli along a circle and makes them orbit around a common centre, thus forming a cluster that moves as a block.

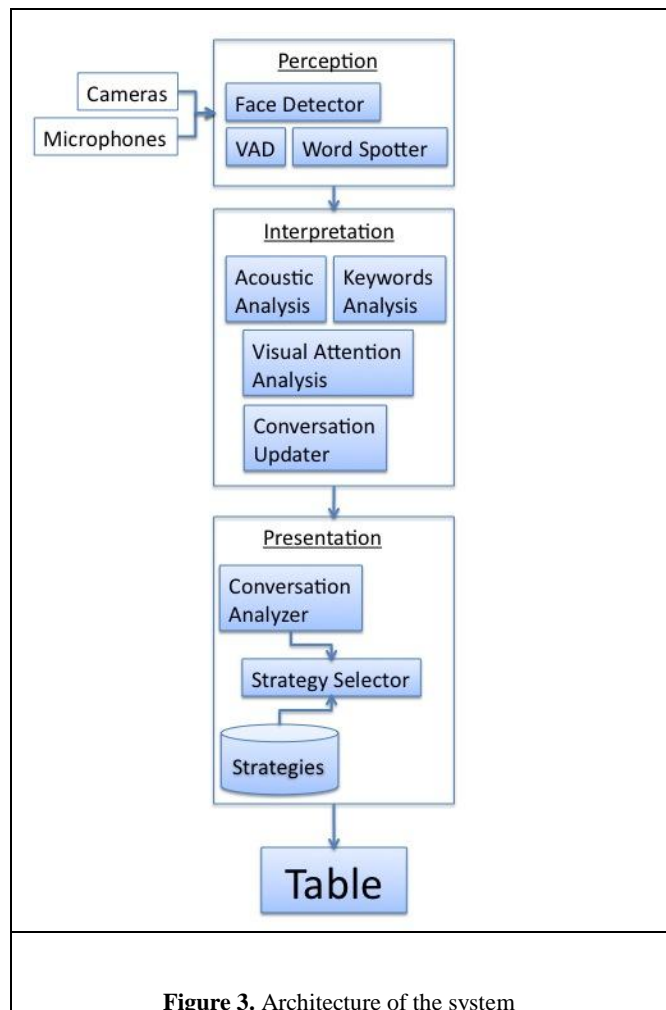


Figure 3. Architecture of the system

The *transition* dimension is related to the notification one. Every action that changes the current visual state of the display can be considered a transition. Transitions exploit graphical effects to attract an appropriate amount of attention from the users and affect the development of the conversation. For example, a topic shift can be suggested by having objects related to the old topic disappear and objects related to the new one

progressively appear; to represent that a topic t is proposed as the prevailing one, the object related to t can be progressively enlarged while the others gradually scale down.

5 PRESENTATION STRATEGIES

We devised a set of presentation strategies to dynamically select and move stimuli. Strategies are organized according to the three dimensions presented above. The selection of strategies is based on the current perceived state. Each strategy has a goal. Goals implemented at the moment are: (i) support the current conversation, (ii) engage a member in a conversation, (iii) start a conversation.

The first strategy is an implementation of Fogg's conditioning strategy. It happens when somebody is speaking and most people are looking at the table. Here we assume that there is an ongoing conversation and the system supports it. Data represented can be either images or words, animations are slow because the notification level is low, transitions are smooth, to not demand too much attention.

The second type of strategy has the goal to engage a "passive" member in a conversation. This strategy is applied when only one's attention is directed to the table, regardless there is an ongoing conversation. The core of the strategy is to show a stimulus directly toward the "passive" member, in order to suggest a possible topic to discuss. This strategy selects only one image to be shown, for it is more suggestive than a word. The notification level is medium because her attention is already directed toward the table and the goal is to make her aware of a possible topic to be discussed. Transition is medium, because the user has to notice the difference with respect to the previous state of the tabletop.

The third strategy, as the second, is an implementation of Fogg's suggestion strategy. It is applied when the conversation stalls and the goal is to engage the group to discuss a topic related to the visit. It is realized by showing an enlarged image, possibly with many details, which represent possible topics to discuss. The notification of this strategy is very high, for people's attention has to be directed toward the table. The transition level is also high, to ensure that at least somebody notice the change.

Finally, a different type of strategy is based on the idea of using humor to trigger interest. It consists of displaying a fun verbal expression followed, after a short time by an image that refers to it in some way. Like in many broadcast ads we see today on newspapers, wall or TV, slight variations on well known linguistic expressions surprise the audience and get its attention. Normally it is a form of irony that plays on the substitution of an element in the expression with a word that evokes the concept that the ad intends to promote. An example for our case is "Saturday knight fever", a variation on the well known movie title that evokes the scene of the festive tournament, part of one of the paintings in the museum we are conducting the experiments in. For this study, the humorous expressions were compiled by hand, but see [10] for a system that automatically produces such puns taking into account the context.

6 USER STUDIES

We conducted an observational study on a prototype which implements a subset of the features presented above. As said, the current system works on the scenario of a visit at the "Cycle of the Months" frescoes. This artwork consists of eleven side-by-side frescos each one measuring on average 2 meters wide and 3 meters high, and representing a particular calendar month. The frescos were painted during the 15th century and illustrate the activities of aristocrats and peasants throughout a full year. The main topics for conversation are therefore the eleven frescos (named after the months they represent) and the keywords that can be recognized by the word spotter are 50 words for details and objects depicted in the frescos. The image repository includes images of the full frescos and of relevant details.

For practical purposes, the studies were run at our labs, using two rooms: in the first the Torre Aquila exhibition was partially reproduced; the second room was used as the post-visit meeting place and it was equipped with a table together with top projection, and a camera to record the sessions.

The subjects in a group of 3 or 4 were initially welcomed and asked to visit the reconstruction of the frescoed room. Each of them received a booklet describing the frescoes. To induce a controlled difference in the experiences, two subjects received a booklet slightly different from the others: one contained more details about some of the frescos and the other some information about the restoration process. In order to reduce the bias and to avoid getting too much attention to the table since the beginning of the experiment, the subjects were told that the purpose of the study was to assess the quality of the information provided in the booklet. After the visit, the subjects were accompanied in the other room, where the only available piece of furniture was the system table, and asked to wait for the experimenter to come back. Soon after, the table started to display stimuli according to the phases described above. After approximately ten minutes the experimenter came back and the subjects were debriefed with a short unstructured interview about their experience with the table; the topic of the interview were people's feelings and attitudes towards the table, the way it functioned, and its place in a real museum.

A total of 5 groups participated in the study; subjects were balanced with respect to gender. Their ages range from 35 to 45. They were all volunteers and, with the exception of two computer programmers, the others had no specific technical skills. Few of them had been in Torre Aquila. All the sessions were video recorded.

The following discussion is based on qualitative observations of the videotapes of the interactions and on the unstructured interviews.

6.1 DISCUSSION

In general the table triggers some interest and it mainly fosters conversation about the technology itself and only indirectly it supports reflection on the actual visit. Yet, in several cases, the appearance of an image leads to a discussion about the fresco, and in particular, for those details for which the booklet does not provide enough information. Mostly, this persuasive effect of the system takes place in moments when the group is temporarily silent. The table is looked at with more attention, and stimuli for continuing the conversation are sought. When the conversation

revives again, usually the table and the stimuli subsequently displayed are ignored. This may be due to the fact that when the table itself is not the topic of the discussion, it is a relevant source of inspiration only when the stimuli are needed to overcome a moment of social lack of interaction; but when the group can independently sustain the conversation, further stimuli are not needed and the group tends to ignore the table.

This may lead to conclude that the phase (a) above is more useful than phase (b). Yet, a possible source of bias is due to the low performance of the word spotter when the conversation is lively because of the frequent overlaps in speech of the different individuals. Therefore, many times the stimuli presented in phase (b) were not actually related to the topic of the conversation but they appeared rather randomly selected.

Another possible bias to the interpretation that the table is not effective in supporting a conversation already in place is that our system does not consider the visual attention of the individuals while proposing the stimuli. While this is less problematic when the conversation is lagging (since the individual tries to get inspiration from the table), it is more important when the conversation is lively since the social pressure tends to focus the individual attention toward the group rather than on the table. A system able to monitor the visual attention of the individuals may choose the right moment to display a stimulus and possibly make phase (b) more effective.

Regarding the use of humorous sequences of text snippets and related details, we failed to observe any interesting effect on the group. Yet this is mainly due to the fact that the sequences went always unnoticed by the subjects. The witty expressions had no special characters and floated similarly to the captions. Again, it seems that the main reason is the lack of consideration about the visual attention when this strategy is applied. This is more problematic with respect to the case of the simpler visual strategies because attention must be assured from the beginning to the end of the sequence in order to get the humorous effect.

Regarding the possibility of interaction with the table, contrary to the subjects of other WoZ studies [9], in general our subjects did not expect the table to be interactive (though someone mentioned the possibility of using Microsoft Surface). Yet, when asked, they expressed a positive attitude toward the idea of using the table as an interactive kiosk and also suggested other activities, as for example accessing email and news.

7 CONCLUSIONS

We can say that the current prototype has a moderate effect on the conversation. With the current implementation, the effect is more prominent for phase (a), when the group conversation is out of topic and in particular during the moments of silence when the individuals are looking for new topics of discussion. To better explore phase (b), we need to refine the module that monitors the conversation and in particular the word spotter that has a reduced accuracy when the conversation is more lively, so as to have a more accurate assessment of the topic being discussed.

REFERENCES

[1] G. Leinhardt and K. Knutson, *Listening in on Museum Conversations*. Altamira Press, 2004.

[2] B. J. Fogg, *Persuasive Technology: Using Computers to Change What We Think and Do* Morgan-Kaufmann, 2002.

[3] C. Rocchi, O. Stock, M. Zancanaro, M. Kruppa and A. Krüger, *The Museum Visit: Generating Seamless Personalized Presentations on Multiple Devices*. Proceedings of the International Conference on Intelligent User Interfaces, IUI-2004, Island of Madeira, Portugal, 2004

[4] D. Petrelli and E. Not, 'User-centred Design of Flexible Hypermedia for a Mobile Guide: Reflections on the HyperAudio Experience' *User Modeling and User-Adapted Interaction*, vol. 15, numbers 3-4, pp. 303-338 (2005)

[5] J. Wilpon, L. Rabiner, C. Lee and E. Goldman, 'Automatic recognition of keywords in unconstrained speech using hidden Markov models', *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, pp. 1870-1878, 1990.

[6] M. Weiser, M. and J. S. Brown, 'Designing Calm Technology', *PowerGrid Journal*, v1.01, July 1996

[7] T. Matthews, A. K. Dey, J. Mankoff, S. Carter and T. Rattenbury, *A Toolkit for Managing User Attention in Peripheral Displays*. Proceedings of UIST '04, Santa Fe, NM, 2004.

[8] H. Ishii and B. Ullmer, *Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms*, Proceedings of Conference on Human Factors in Computing Systems CHI. Atlanta, GA, pp. 234-241, 1997.

[9] J. F. McCarthy, T. J. Costa and E. S. Liongosari, *UniCast, OutCast & GroupCast: Three Steps Toward Ubiquitous, Peripheral Displays*, Proceedings of the 3rd international Conference on Ubiquitous Computing, Atlanta, GA, 2001

[10] W. W. Gaver, J. Bowers, A. Boucher, H. Gellerson, S. Pennington, A. Schmidt, A. Steed, N. Villars and B. Walker, *The drift table: designing for ludic engagement*. CHI '04 Extended Abstracts on Human Factors in Computing Systems, 2004.

[11] N. A. Streitz, C. Rocker, T. Prante, D. V. Alphen, R. Stenzel and C. Magerkurth, 'Designing Smart Artifacts for Smart Environments', *Computer*, vol. 38, no. 3, pp. 41-49, 2005.

[12] J. M. DiMicco and W. Bender, *Group Reactions to Visual Feedback Tools* Proceedings of The Second International Conference on Persuasive Technology, Stanford University, CA, 2007.

[13] J. Sturm, O. Houben-van Herwijnen, A. Eyck, and J. Terken *Influencing Social Dynamics in Meetings through a Peripheral Display*, Proceedings of International Conference on Multimodal Interfaces, ICMI 2007, pp. 263-270, 2007.

[14] T.J. Kim, A. Chang, L. Holland and A.S. Pentland, *Meeting Mediator: Enhancing Group Collaboration with Sociometric Feedback*, Proceedings of ACM Conference on Computer-Human Interaction (CHI08), Florence, 2008.

[15] D. Boud, R., Keogh, and D. Walker (Eds.), *Reflection: Turning Experience into Learning*. Kogan Page, 1985.

[16] G. Adomavicius, and A., Tuzhilin, 'Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions', *IEEE Transactions on Knowledge and Data Engineering* 17 (6), pp. 734-749, 2005.

[17] C. Rocchi, D. Tomasini, O. Stock and M. Zancanaro, *Fostering conversation after the museum visit: a WOZ study for a shared interface*. Proceedings of the Conference on Advanced Visual Interfaces, AVI-2008, Napoli, Italy, 2008

[18] P. Viola and M. Jones, 'Robust real-time face detection', *Computer Vision*, 57(2) pp.137-154, 2004.

[19] L. Bartram, C. Ware and T. Calvert, 'Moticons: Detection, distraction and task', *International Journal Human-Computer Studies*, 58 (5). pp. 515-545, 2005.

[20] R. Abrams and S. E. Christ, 'Motion onset captures attention: A rejoinder to Franconeri and Simons (2005)', *Perception and Psychophysics*, 68 (1), pp.114-117, 2006.