On the Need for a Framework for Attentive Groupware Systems

António Ferreira and Pedro Antunes

Department of Informatics, University of Lisbon, Portugal
asfe@di.fc.ul.pt, paa@di.fc.ul.pt

Abstract In this paper we present a combined perspective over groupware systems and attentive user interfaces to address the balance between the rising needs for collaboration with the rising costs of not paying attention and being interrupted. We introduce a framework that derives attentive devices from groupware information flows to take into account attentional phenomena. Finally, we demonstrate the framework's role in an attentive brainstorming tool and provide preliminary results from a pilot experiment.

1 Motivation

A continuing trend in groupware research aims at improving the sense of proximity within groups, whether by enabling geographically distant people to work together, or by supporting ongoing activities at different times. Researchers have been proposing several mechanisms to enhance group awareness by providing *ever greater* information about the presence and actions performed by people, e.g., radar and fish-eye views, multi-user scrollbars, and telepointers [1]. A main argument is that communication channels mediated by computers are relatively poor when compared with more natural settings, such as face-to-face meetings [2]. However, a problem with this trend is that it fails to recognise that sometimes *more is less* due to the limitations of the human attentive capacity, especially as we become surrounded with computers and, not necessarily useful, information [3].

During the late 1990s several researchers from the Human-Computer Interaction (HCI) field have became interested in Attentive User Interfaces (AUI), and since then this topic is gaining momentum, as evidenced by a special issue in the Communications of the ACM [3], and a second one in Computers in Human Behavior [4]. The prime motivation for AUI is the recognition that as the needs for information and communication rise so do the costs of not paying attention and being interrupted [3]. So, instead of assuming the user is always focused on the entire computer display, AUI negotiate the users' attention by establishing priorities for presenting information.

Most research on AUI is directed towards single-user activity, the main assumption being that individual performance degrades with the number of simultaneous requests for attention. Therefore, researchers are enhancing input/output devices so that the user remains focused on a primary task without getting too much distracted by a secondary, typically unrelated, task, e.g., by using eye-trackers [5], statistical models of human interruptibility [6,7], displays that show information at various levels of detail [8], and other devices.

Regarding multi-user activity, the literature is mainly situated in video conferencing [9,10], making the study of attention in computer-mediated collaborative contexts a largely unexplored area. Moreover, the current emphasis is directed towards evaluating the enhanced input/output devices themselves, in contrast with determining the outcomes of using these devices in work settings. Furthermore, the convergence of groupware and AUI poses new challenges because of the differences between individual and group work, such as: a) people working in a group are exposed to more interruptions because they have to manage more information flows; b) instead of doing a single, extensive task, group members usually execute a series of intertwined tasks; and c) in group work the primary and secondary tasks are typically related and can both contribute to the shared goal.

Ferreira, A. and P. Antunes (2007) On the Need for a Framework for Attentive Groupware Systems. 1st Workshop on Adaptation and Personalization in Social Systems: Groups, Teams, Communities, Corfu, Greece.

Given this situation, we propose a framework for mediating attention in groups and put forward the possibility of using specialised, software-based, groupware devices to account for attentional phenomena. These *attentive devices* would work by manipulating the information flows supported by the groupware system, and since, in this work context, collaboration is mediated by computers, its use might mean that the group as a whole would be more attentive, and possibly more productive.

We explain the framework in Sect. 3 and in Sect. 4 we describe its application to an electronic brainstorming tool. In Sect. 5 we conclude the paper with a discussion regarding our approach and with prospects for future work.

2 Related work

The study of the mediating role of computers on attention within groups is largely an unexplored research area, with the exception of video conferencing. Currently, the major part of the literature on AUI is focused upon single-user software (singleware). We will refer to both multi and single-user contexts to provide a more comprehensive picture of the systems and devices that address human attention in HCI.

2.1 Attentive groupware systems

GAZE is a groupware system developed to facilitate the detection of who is talking to whom in remote meetings [11]. It works by displaying photographs of users on the computer display, which can be rotated by intervention of eye-trackers placed in front of each user. For example, the photos might be rotated towards a photo of a user who is speaking. In this way, turn taking may be more natural and require less interruptions to determine who will speak next.

The GAZE-2 system was developed to overcome the technical limitations of the original GAZE and to support multiple conversations at the same time [9]. In GAZE-2 each user has three video cameras that capture the user's face at slightly different angles; then, an automatic camera director chooses one video stream taking into account the direction where the user is looking at. As in the original GAZE, the representation of each user is rotated to reflect the focus of attention. Moreover, as the angle of rotation increases, the quality of the video stream is purposefully reduced to save network bandwidth. According to the authors, this technique is effective because it is based upon our own natural limitations concerning peripheral vision. Another feature of GAZE-2 is the automatic filtering of voices when multiple conversations are being held at the same time. Depending upon the user in focus, so is the respective audio stream amplified, and the other streams attenuated (but not eliminated). If the focus of interest suddenly changes, as sensed by the eye-tracker, the audio is again adjusted.

Recent work with groupware systems further explore the ideas in GAZE-2. For instance, eyeView supports large meetings by manipulating the size of the video windows and the voice volumes of each user on the group as a function of the current focus of attention [10].

2.2 Attentive devices on singleware systems

In contrast with groupware systems, several input/output devices have been tested on attentive interfaces for single-user applications, such as: a) sensors that detect the user's focus of attention based upon eye-gaze and body orientation [5,12,13]; b) physiological sensors that assess the user's mental workload by measuring heart rate variability, pupil dilatation, and eye-blink activity [14,15,16]; c) statistical models that determine adequate moments to interrupt and communicate with the user [6,7]; and d) displays that present information at various levels of detail, depending on the user's focus of attention [8].

Regarding the use of eye-trackers to support human attention, applications include magnifying the graphical window in which the user is currently focused on, controlling a robotic directional microphone coupled to a video camera, to overhear a particular conversation taking place in a remote room, and detecting eye contact to automatically choose which electronic appliance should obey to voice commands [5]. Other applications use eye-gaze to position a cursor on the screen with minimal manual intervention [12] and to help a user read a book written in a foreign language [13].

Body orientation sensors have been tested in an office environment to regulate the transparency of cubicle walls (opaque when the user does not wish to attend requests from others) and to control noise cancellation in headphones [5].

Concerning physiological sensors, these have been used to assess mental workload, which, in turn, is considered a surrogate of the user's attentional state. One study suggests using heart rate variability and electroencephalogram analysis to distinguish between at rest, moving, thinking, and busy states, and describes an automated regulator of notifications that can be installed on mobile phones [15]. Heart rate variability had already been used in the 1990s to assess conditions of excessive mental effort [16]. More recently, pupil dilatation has been used, in combination with hierarchical task models, to predict opportune moments to interrupt the user [14].

Another approach to detect the best time to interrupt the user is to apply statistical models that permanently estimate and balance the value of information with the cost of interrupting, based upon a stream of clues, such as, appointments on the personal calendar, past activities and work patterns, ambient noise, or body posture [6]. Statistical models have also been used to select the best predictors of interruptibility from a myriad of software sensors embedded in an integrated development environment [7].

Finally, we also refer to a special display that adjusts the level of detail in selected areas of the screen as a function of the user's current visual focus of attention [8]. This effect has similarities with the GAZE-2 approach.

2.3 Discussion

Research in AUI, its applications and devices, is progressing in many directions. However, we argue that most studies are directed towards evaluating the devices *per se*, in contrast with determining the outcomes of using these devices in work settings. For example, the GAZE-2 system was evaluated through a user questionnaire that measured the self-perception of eye-contact and distraction, as well as changes in colour and brightness during camera shifts [9], but no attempt was made to determine if group work benefited. The same situation occurs with eyeView [10], which is still in an early stage of development.

Some studies do address the evaluation of AUI from the perspective of task execution, but are restricted to single-user activity. One study measured the effects of interruption on completion time, error rate, annoyance, and anxiety, and suggest that AUI should defer the presentation of peripheral information until task boundaries are reached [17]. In another study, the effectiveness and efficiency of users were evaluated as they performed two types of tasks under the exposure of four methods for coordinating interruption, and recommends that AUI should let users manually negotiate their own state of interruptibility, except when response time for handling interruptions is critical [18].

We note, however, that there are numerous differences in individual and group work that might reduce the external validity of current results. This opens an opportunity for doing research in groupware systems and AUI.

3 Framework for attentive groupware systems

The purpose of this framework is to conceptualise attentive groupware systems. Its underlying assumption is that group performance may improve by incorporating human attentional phenomena in attentive devices that adjust the groupware mediation (e.g., the information in circulation) as people carry out the collaborative tasks (see Figure 1).

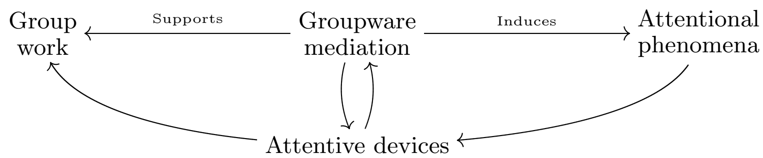

Figure 1: Conceptual framework for attentive groupware systems

Performance evaluation of group work is outside of the scope of the framework but is strongly tied to its usefulness. In the next sections we look into the concepts in the framework.

3.1 Groupware mediation

Conceptually, the groupware system is at the centre of the group. It is a mediator that deals with all sorts of information that comes and goes between and among the users. Moreover, we assume that the users are restricted to using groupware to collaborate, e.g., when they are geographically distributed, since this is a challenging scenario for managing attention in groups.

Groupware mediation should support the notions of interdependence, e.g., for group planning, and mutual awareness, e.g., for situation assessment, which generate information that requires human attention in order for the group to make progress. Our strategy to characterise groupware mediation is to change the perspective over the information flows that we investigated in a previous study about shared workspaces [19]: instead of looking into the ways remote people collaborate, we analyse the corresponding mediation modes.

In the *explicit communication mode* the groupware receives information produced by a user and forwards it to one or more users, based upon an explicit request by the sender [2]. This may happen, for instance, when a user expresses a request for an object to the user who is holding it; another example is when an instructor provides online guidance to students for collaborative problem solving (illustrated in Figure 2a[1]). This information mode can be supported by a groupware interface capable of multiplexing information from input devices to several output devices, e.g., a user typing on a keyboard and the other users seeing the text on their displays.

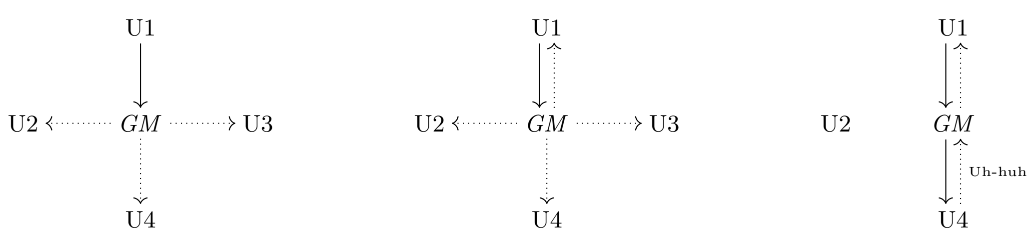

(a) Explicit communication (b) Feedthrough (c) Back-channel feedback

Figure 2: Groupware mediation modes

In the *feedthrough mode* the groupware automatically reports actions executed by one user to several users [20]. This mode is essential because it provides group awareness and enables users to construct meaningful contexts for collaboration. For example, a graphical shared workspace may provide its users with information about the menu selections for each user who is manipulating objects. The groupware interface can support the feedthrough mode by capturing each user's inputs and then multiplexing feedback information (replies to a single user) to the other

[1] In Figure 2, *GM* means groupware mediation and Un is user n.

users on the group (see Figure 2b). Sophisticated schemes may consider delivering less information by manipulating the granularity and timing associated with the operations executed through the groupware [21]. Interestingly, the motivation for these schemes has been related to technical limitations, such as network bandwidth, but, in our view, the limitations of the human attentive capacity should also be accountable because the amount of information generated by the groupware may overcome us, and thus decrease group performance.

In the *back-channel feedback mode* the groupware captures unintentional information initiated by a user and directs it to another user to facilitate communication and to convey human states of attention. This may occur, for instance, when a listener says 'uh-huh' to indicate that s/he is following the speaker (see Figure 2c). To capture this type of information the groupware interface can use attentive devices such as those described in Section 2, or use other sensors that take advantage of the information that is available to the groupware.

One of the concepts in the framework embraces specialised attentive groupware devices, but first we look into some phenomena related to human attention.

3.2 Attentional phenomena

Human attention is often associated with the selection of relevant information and simultaneous attenuation or discard of non-relevant data. It is a process that optimises the use of our limited cognitive resources so that we can perceive or act accurately and quickly [22,23,24,25].

Attention and consciousness are though to be different processes: whereas attention covers the full spectrum of data that we manipulate (through our senses or internal memory), consciousness is confined to the information that we are currently aware off [24]. Furthermore, a recent study identifies scenarios in which attention might not give rise to consciousness and vice-versa [22]. This means, for instance, that we may be scanning a computer display without noticing important information.

Over the decades, psychologists and cognitive scientists have been identifying the goals and limitations of human attention. Researchers should pick up this knowledge to invent and evaluate ways to support the goals and compensate the limitations.

Two of the main goals of attention are accuracy, to perceive specific objects and to execute particular tasks, and speed responding, to perceive objects or execute tasks after the presentation of a predictive cue [26]. Accuracy manifests itself when we successfully remove or attenuate the influence of extraneous and confusing information. The 'cocktail party' phenomenon—our ability to keep track of a conversation in a crowded room—is related to this goal [24,23]. Speed responding occurs when we are able to respond faster to anticipated events [26] and almost always involves a clear expectation of when to initiate the response [23].

An example of a groupware system that addresses attentional accuracy is GAZE-2 [9], which automatically regulates the sound volume of overlapping conversations according to the user's current visual focus of attention. The effects on group performance have not, however, been evaluated, and more contexts of group work, besides remote meetings, may benefit from explicit support of human attentional goals. Regarding groupware support to speed responding, anticipation of upcoming events is actually quite possible in computer-mediated work because, as soon as the system detects user activity that may be pertinent to another user, it can signal that activity and also control the delivery time to create clearer expectations. For example, popular instant messengers provide anticipatory cues when users start typing a message.

Interestingly, human attention has its own limitations, in that it can fail to select the relevant information, or take too long. Moreover, some phenomena are known to occur even after tremendous training, such as the 'attentional blink,' which is a delay between paying attention to one object or task and attending to the next stimulus [25]. On the other hand, there is evidence that the response time to the second stimulus may be reduced if the time between attention switches is made longer and, in particular, constant [27]. This type of intervention should be performed by groupware systems to regulate the flows of information that each user is exposed to.

Another attentional phenomenon is 'change blindness,' which manifests itself when we fail to notice changes, even dramatic ones, such as a swap of the person with whom we were talking to just seconds ago. As long as the change matches the context, e.g., swapping of students during a brief encounter in a university campus, we may simply miss the difference [23]. If we do want to check if anything has changed, then we may have to engage in a very slow process of scanning the full picture before us. This happens because, although we can attend to four or five objects simultaneously, we can only detect one change at a time [28]. The consequences of 'change blindness' in people doing group work should be apparent. The notion of group lends itself to the creation of a social context and the existence of several people collaborating, i.e., contributing to the same shared goal, stimulates scenarios in which multiple changes may occur simultaneously. This creates the required conditions for people not noticing changes, which may reduce group performance because of the time needed to catch up. Groupware system should, therefore, highlight changes to compensate this attentional limitation.

3.3 Attentive devices

The last concept in the framework comprehends the input/output devices that support the groupware mediation. We propose a classification that comprises awareness and coupling devices, which may themselves be manipulated by specialised attentive devices.

We define *awareness input/output devices* as devices dedicated to sensing and displaying information about the collaborative activity within the group, allowing users to construct a perceptual image of the work context. This information is usually designated 'group awareness' [2]. A myriad of these devices have been described in the literature [1,2], and indeed, we argued at the beginning of the article that groupware research has been focusing on these devices.

The *coupling input devices* are used to loose the link between the actions executed be a user and the information that is passed on to the other users [29]. Two types of coupling control may be considered in groupware mediation: first, coupling may be exerted at the origin to specify what and when information produced or about a user should become available to the rest of the group; second, coupling may be applied at the destination to apply filters that restrict group awareness to some selected objects and actions.

Figure 3 illustrates that users may use their respective coupling input devices to control the outflow of information from both the input and output awareness devices. Note that coupling control does not apply to (single-user) feedback.

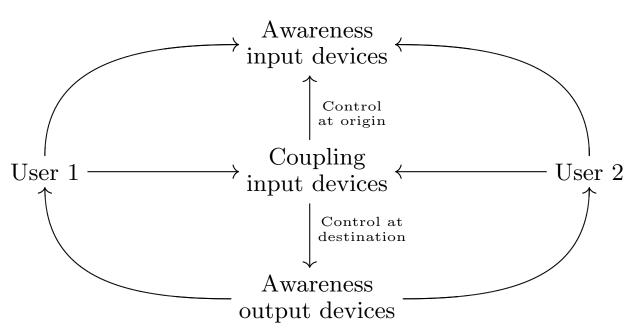

Figure 3: Groupware input/output devices

Coupling devices require manual discrimination and control of group awareness, thus penalising individual performance. However, this disadvantage is balanced by the capacity to limit the amount of information, which may improve attention within the group. This tradeoff sets the stage for introducing specialised *attentive devices* for groupware systems.

We propose a set of attentive devices that collect and combine information received from awareness sensors associated with each user, and that automatically manage the information that is delivered to the awareness displays, according to human attentional phenomena:

Activity anticipator (AA) Senses users' actions, or lack of activity, that may affect group performance and delivers preliminary information (cues), to prepare users to be attentive to upcoming outcomes and to enable faster response times.

Change emphasiser (CE) Tracks awareness information available on each user's display and highlights changes caused by the latest group activities, to attenuate the effects of the 'change blindness' phenomena and to help users make faster situation assessments.

Opportunity seeker (OS) Senses activity on a user basis and seeks for opportune moments to deliver group awareness, such as after the completion of a subtask, to enable faster response times to stimulus.

Time separator (TS) Intercepts the delivery of group awareness to the users, and introduces a constant delay, after which the hand over proceeds. The delay should be constant from the point of view of each user, and its purpose is to attenuate the effects of the 'attentional blink' phenomena and to improve task switching performance.

Some of these attentive devices may be combined to achieve more complex manipulation of groupware mediation. For example, TS → CE has the effect of highlighting changes based upon group awareness collected over a constant period of time, instead of as soon as the changes occur. However, some combinations may be incompatible; for instance, TS → OS are in conflict because they incorporate different expectations about when to deliver information to the users.

This set of attentive devices, which may incorporate more devices in the future, should be applicable to groupware mediation in a broad range of collaborative scenarios. For example, in asynchronous groupware, the change emphasiser may be used to highlight differences between two discrete states of group work. In the next section we describe an application of these four attentive devices to an electronic brainstorming tool.

4 An attentive brainstorming tool

We applied the proposed framework to the design and development of an electronic brainstorming tool, in the context of a laboratory experiment for testing the effects on group performance caused by the attentive devices.

4.1 Architecture and components

The attentive brainstorming tool is characterised by a client-server architecture, in which groupware mediation is performed on the server. The purpose of the clients, one per user, is to receive input from the users and pass it on to the server, to display group awareness as it becomes available from the server, and also to collect performance data that is stored in a server log.

Groupware mediation is entirely supported in the feedthrough mode, i.e., the brainstorming tool automatically distributes new ideas to all users, without requiring mechanisms for explicit or back-channel communication between or among the users. All users remain anonymous during the brainstorming sessions.

In this work context, group performance depends on the number of ideas that the users put forward, and this, in turn, is related to the quantity of ideas that each user is able to attend to, in particular, ideas coming from the other users. The purpose of the attentive devices is to enable users to attend to more ideas, which may have a positive net impact on group performance. The brainstorming tool implements all four attentive devices: a) the activity anticipator senses group inactivity and alerts brainstormers that the session will soon end unless new ideas are put forward; it also displays a message indicating that a user is typing an idea (see bottom zone in

Figure 4); b) the change emphasiser highlights ideas that have been recently put forward by other users (see horizontal line separating old and new ideas in Figure 4); c) the opportunity seeker waits for a user to stop typing to then display new ideas; and d) the time separator shows new ideas at constant time intervals. The attentive devices may be individually enabled or disabled.

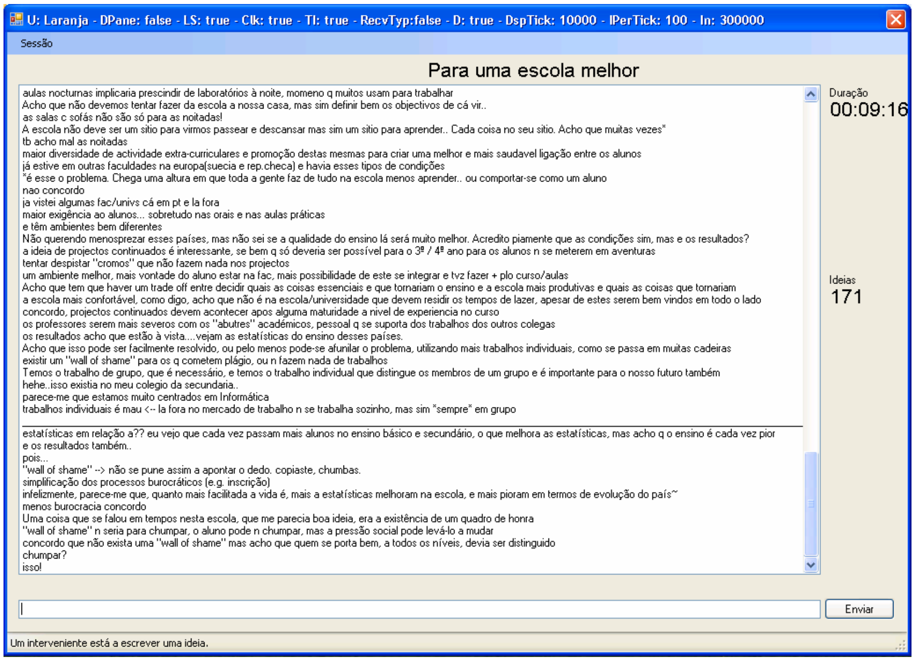

Figure 4: Attentive brainstorming tool

4.2 Preliminary results

We now briefly describe a pilot experiment that we conducted in an university laboratory with a group of five male graduate students, gathered for a brainstorming session through the attentive tool. No other means of communication, such as speaking or gesturing were allowed, to simulate remote collaboration. The session lasted for thirty five minutes and the topic was 'for a better school.' Our expectation was that this topic would induce many ideas, so as to create a scenario that demanded more attention from the users. The ideas did came, indeed, but at the cost of breaking the rules of brainstorming because the session was transfigured into a chat (as illustrated in Figure 4; the average number of characters per idea was seventy two). We decided to continue the experiment for three reasons: a) to observe the users' activity under the exposure of the attentive devices; b) to check if the attentive devices were working as expected; and c) to obtain performance data for posterior analysis.

We did not notice any unusual behaviour from the users, such as apprehension or frustration, and adaptation to the tool required no training. One of the attentive devices malfunctioned: the change emphasiser was not highlighting new ideas as expected, and we latter found out that the cause was linked to faulty interactions with the time separator. There was also a minor glitch with the opportunity seeker in that between one and seven ideas slipped through to the display while users where typing an idea, which was not supposed to happen. We also detected that several key-press events in the performance log were recorded with the same time stamp for the same user, which probably indicates lack of timing precision.

Table 1 shows a summary of the performance data that were collected by the brainstorming tool. We randomly assigned colours to the five users, but unfortunately two of them were given the same colour during the configuration of the tool, so we exhibit only results for three users.

Table 1: Results from the pilot experiment

	Produced ideas			Received ideas		
User	#	Typing	s/idea	#	Processing	s/idea
Green	65	17m 09s	15.3	240	6m 03s	1.5
Red	42	10m 40s	9.3	234	6m 40s	1.7
Orange	68	20m 49s	18.4	244	5m 05s	1.2

The total number of ideas contributed during the session, for the *five* users, was 244. User 'orange' received all of the ideas, but the other two users received less because the session ended before the time separator could deliver the last remaining ideas. This should cause no problem in practice. The rightmost column in Table 1 shows the number of seconds of non-typing activity immediately after the reception of ideas, divided by the total number of ideas that the user received throughout the entire session. If we sum the values in this column, on a per user basis, with the average number of seconds to type an idea, then, in future experiments, we are expecting up to 10–20 seconds for situations in which the user types a contribution in reaction to an idea that s/he received from another user. We also note that during the thirty five minutes that this session lasted, between 30–60% of the time, per user, was spent on typing ideas.

5 Conclusion and future work

The study of the mediating role of computers on human attention in remote collaboration settings is largely an unexplored research area. While current trends keep aiming at conveying ever greater information about the group, we suggest a route that explicitly recognises the limitations of the human, and therefore the group, attentive capacity. This route is consistent with AUI research—which is almost entirely directed at single-user activity—and we argue that the existing body of knowledge should extended into the groupware field. To this end we introduce a framework for attentive groupware systems and hypothesise that group performance improves with the use of specialised attentive devices.

Many questions remain unanswered: does group performance significantly improve? What attentive devices are best suited for different collaborative scenarios? Can groups be made larger while remaining attentive and productive? We are currently addressing some of these questions with an attentive brainstorming tool, and setting up experiments with groups of five people. The road is open for many more experiments and applications.

References

1. Raikundalia, G.K., Zhang, H.L.: Newly-discovered group awareness mechanisms for supporting real-time collaborative authoring. In: AUIC'05: Proceedings of the sixth Australasian conference on User interface, Newcastle, Australia (2005) 127–136
2. Gutwin, C., Greenberg, S.: A descriptive framework of workspace awareness for real-time groupware. Computer Supported Cooperative Work **11**(3) (2002) 411–446
3. Vertegaal, R.: Attentive user interfaces. Communications of the ACM **46**(3) (2003) 30–33
4. Roda, C., Thomas, J.: Attention aware systems: Introduction to special issue. Computers in Human Behavior **22**(4) (2006) 555–556
5. Vertegaal, R., Shell, J.S., Chen, D., Mamuji, A.: Designing for augmented attention: Towards a framework for attentive user interfaces. Computers in Human Behavior **22**(4) (2006) 771–789

6. Horvitz, E., Kadie, C., Paek, T., Hovel, D.: Models of attention in computing and communication: From principle to applications. Communications of the ACM **46**(3) (2003) 52–59
7. Fogarty, J., Ko, A.J., Aung, H.H., Golden, E., Tang, K.P., Hudson, S.E.: Examining task engagement in sensor-based statistical models of human interity. In: CHI'05: Proceedings of the SIGCHI conference on Human factors in computing systems, Portland, Oregon, USA (2005) 331–340
8. Baudisch, P., DeCarlo, D., Duchowski, A.T., Geisler, W.S.: Focusing on the essential: Considering attention in display design. Communications of the ACM **46**(3) (2003) 60–66
9. Vertegaal, R., Weevers, I., Sohn, C., Cheung, C.: GAZE-2: Conveying eye contact in group video conferencing using eye-controlled camera direction. In: CHI'03: Proceedings of the SIGCHI conference on Human factors in computing systems, Ft. Lauderdale, Florida, USA (2003) 521–528
10. Jenkin, T., McGeachie, J., Fono, D., Vertegaal, R.: eyeView: Focus+context views for large group video conferences. In: CHI'05: Extended abstracts on Human factors in computing systems, Portland, Oregon, USA (2005) 1497–1500
11. Vertegaal, R.: The GAZE groupware system: Mediating joint attention in multiparty communication and collaboration. In: CHI'99: Proceedings of the SIGCHI conference on Human factors in computing systems, Pittsburgh, Pennsylvania, United States (1999) 294–301
12. Zhai, S.: What's in the eyes for attentive input. Communications of the ACM **46**(3) (2003) 34–39
13. Hyrskykari, A., Majaranta, P., Aaltonen, A., Räihä, K.J.: Design issues of iDICT: A gaze-assisted translation aid. In: ETRA'00: Proceedings of the 2000 symposium on Eye tracking research and applications, Palm Beach Gardens, Florida, USA (2000) 9–14
14. Iqbal, S.T., Adamczyk, P.D., Zheng, X.S., Bailey, B.P.: Towards an index of opportunity: Understanding changes in mental workload during task execution. In: CHI'05: Proceedings of the SIGCHI conference on Human factors in computing systems, Portland, Oregon, USA (2005) 311–320
15. Chen, D., Vertegaal, R.: Using mental load for managing interruptions in physiologically attentive user interfaces. In: CHI'04: Extended abstracts on Human factors in computing systems, Vienna, Austria (2004) 1513–1516
16. Rowe, D.W., Sibert, J., Irwin, D.: Heart rate variability: Indicator of user state as an aid to human-computer interaction. In: CHI'98: Proceedings of the SIGCHI conference on Human factors icomputing systems, Los Angeles, California, USA (1998) 480–487
17. Bailey, B.P., Konstan, J.A.: On the need for attention-aware systems: Measuring effects of interruption on task performance, error rate, and affective state. Computers in Human Behavior **22**(4) (2006) 685–708
18. McFarlane, D.C.: Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. Human-Computer Interaction **17**(1) (2002) 63–139
19. Antunes, P., Ferreira, A., Pino, J.A.: Analyzing shared workspaces design with human-performance models. In: CRIWG'06: Proceedings of the twelfth international workshop on Groupware, Medina del Campo, Spain (2006) 62–77
20. Hill, J., Gutwin, C.: Awareness support in a groupware widget toolkit. In: GROUP'03: Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work, Sanibel Island, Florida, USA (2003) 258–267
21. Gutwin, C., Greenberg, S.: The effects of workspace awareness support on the usability of real-time distributed groupware. ACM Transactions on Computer-Human Interaction **6**(3) (1999) 243–281
22. Koch, C., Tsuchiya, N.: Attention and consciousness: Two distinct brain processes. Trends in Cognitive Sciences **11**(1) (2007) 16–22
23. Anderson, J. In: Cognitive Psychology and its Implications. Sixth edn. Worth (2005) 63–105
24. Sternberg, R.J. In: Cognitive Psychology. Third edn. Wadsworth (2003) 65–107
25. Eysenck, M.W., Keane, M.T. In: Cognitive psychology: A student's handbook. Fourth edn. Psychology Press (2000) 119–150
26. LaBerge, D.: Attention. In Bly, B.M., Rumelhart, D.E., eds.: Cognitive Science. Academic Press, San Diego, CA, USA (1999) 44–98
27. Roda, C., Thomas, J.: Attention aware systems: Theories, applications, and research agenda. Computers in Human Behavior **22**(4) (2006) 557–587
28. Simons, D.J., Rensink, R.A.: Change blindness: Past, present, and future. Trends in Cognitive Sciences **9**(1) (2005) 16–20
29. Dewan, P., Choudhary, R.: Coupling the user interfaces of a multiuser program. ACM Transactions on Computer-Human Interaction **2**(1) (1995) 1–39