Attention support in digital environments
Nine questions to be addressed

Claudia Roda
American University of Paris
147, rue de Grenelle 75007 Paris
croda@aup.edu

Abstract

Research in *attention aware systems*, i.e. systems that support users in their attentional choices, promises to address many of the problems related to information overload, cyber collaboration, and mobility, by providing features helping users in coping with attentional limitations. However, the design of attention aware systems necessitates a deep understanding of human attentional processes, of the knowledge a system needs to support those processes, and of the manner in which such knowledge may be acquired. Because the conceptualization of such systems requires understanding users' cognitive states in terms of their past interactions with the environment, the interactivist framework may provide a strong basis for analyzing attention aware systems.

This paper briefly introduces the services that attention aware systems may provide, suggests how *attention* may be modeled within the interactivist framework, and proposes nine questions that may be answered within such framework to gain the knowledge necessary for the creation of attention aware systems.

Keywords: attention, attention aware systems, interactivism

PsycINFO classification: 4010, 2346, 4120

1 Introduction

As networked technologies increasingly provide users with access to information and people anywhere/anytime, the need for systems capable of supporting users in their cognitive resource allocation choices has been increasingly recognized (McCrickard, Czerwinski, & Bartram, 2003; Roda & Thomas, 2006b; Vertegaal, 2003). The design of systems capable of adapting to the needs of individual users operating in specific contexts and under particular constraints is attracting significant interest both in industry and the academy. In particular, research in *attention aware systems*, i.e. systems that support users in their attentional choices (Roda & Thomas, 2006a), promises to address many of the problems related to information overload, cyber collaboration, and mobility, by providing features helping users in coping with attentional limitations. Such features may include intelligent notification and management of interruptions, support for task resumption and multitasking, and awareness mechanisms for collaborative attention allocation. In general, attention aware systems may require a complete rethinking of the metaphors underlying device interface conceptualization (Kaptelinin & Czerwinski, 2007; Roda, Stojanov, & Clauzel, 2006) and of the way we design systems as a whole (see for example the ubiquitous computing (Weiser & Brown, 1996) or ambient intelligence (Kunz, October 2001) literature). At the conceptualization level, attention aware systems will require a much better understanding of how human attentional processes work, and how they might be supported.

This paper introduces the services that attention aware systems aim at providing (section 2), it proposes an interpretation of attention within the interactivist framework (section 3), and formulates nine questions about attention that may be answered by relating the users' mental states to sequences of past interactions with the environment (section 4). The overall objective of the paper is to highlight the gaps still existing between research in cognitive psychology on the subject of attention, and the application of this research in computer science in order to build systems that are respectful of humans' limited cognitive abilities.

2 How can systems support attentional processes?

The services presented here (see (Roda, 2006) for a more complete list of services) illustrate various aspects of attention support in modern, community-based, and information-rich environments (for example, a study reports that at any given moment a user has on average eight windows opened, and spends about 11 minutes on a given task before being interrupted (Mark, Gonzalez, & Harris, 2005)). In such environments the users' activity is characterized by frequent interruptions and multi-tasking, requiring that users explicitly evaluate their attention allocation strategies.

We will concentrate on three types of services:

- Interruption management services are probably the most discussed in the literature and aim at minimizing disruption to the user caused by frequent interruptions.
- Support for task switching and task reminder services endeavor to help users in managing multiple, often interleaved, tasks.
- Self and community awareness services provide the support necessary for users to assess and reflect upon their own and the community's attention allocation strategies.

2.1 Interruptions Management

Although interruptions may bring information to one's attention of potential use in the primary (current) task, or even, in the case of simple primary tasks, facilitate task performance (Speier, Vessey, & Valacich, 2003); it has been widely reported that interruptions increase the load on attention and memory (Gillie & Broadbent, 1989), may generate stress (Bailey, Konstan, &

Carlis, 2001; Zijlstra, Roe, Leonova, & Krediet, 1999), and compromise the performance of the primary task (Franke, Daniels, & McFarlane, 2002; McFarlane & Latorella, 2002; Nagata, 2003; Speier, Vessey, & Valacich, 2003) especially when the user is working on handheld devices in mobile environments (Nagata, 2003). In order to minimize disruption while ensuring that pertinent content is appropriately attended to, the system must make a decision about the *relevance* to the user of the newly available information in the current context, and consequently select notification *timing* and *modality*.

2.1.1 Assessing the relevance of newly available information

Research on interruption management is quite extensive (see for example (Burmistrov, 2005)), however the strategies for evaluating the relevance of newly available information are normally based on static structures rather than being dynamically derived from the observation of users' interaction with system and environment. We seem to be able to develop strategies to guide relevance evaluation and refine them through experience, but the learning processes and the factors underlying the development of these strategies are still largely unknown.

2.1.2 Timing of interruptions

Several studies have demonstrated that the timing of an interruption may make a significant difference in both how easily the information presented is acquired by the user, and on how much disruption the interruption imposes on the task being interrupted (Adamczyk & Bailey, 2004; Czerwinski, Cutrell, & Horvitz, 2000). The solutions for the selection of interruption timing proposed so far are either based on task-knowledge or on sensory-input.

Task-knowledge based timing relies on the analysis of the structure of the task being performed. For example, Bailey and his colleagues (Bailey, Adamczyk, Chang, & Chilson, 2006; Bailey & Konstan, 2006) represent tasks as two-level hierarchies composed of coarse events further split into fine events, and demonstrate that interruptions are less disruptive when presented at coarse breakpoints, corresponding to the completion of coarse events.

Sensory-input-based timing relies on sensor input about user activity to detect best times for interruption. On the basis of the observation that human beings can very efficiently, and given only a very small number of cues, evaluate other's interruptibility, Hudson et al. (2003) propose that interruptibility evaluation is attainable with simple sensors, and that speech detectors are the most promising such sensors. Chen and Vertegaal (2004) instead use more sophisticated physiological cues (Heart Rate Variability – HRV, and electroencephalogram - EEG) to distinguish between four attentional states of the user: at rest, moving, thinking, and busy.

In the Atgentive project (Atgentive, 2005-2007; Roda & Nabeth, 2006) task-knowledge-based and sensory-input-based approaches are integrated by combining knowledge of a detailed task structure (Laukkanen, Roda, & Molenaar, 2007) with simple sensory-input to evaluate the strength of breakpoints for possible interruptions.

2.1.3 Interruption modality

Notification modality may influence an interruption’s impact on user activity at various levels: interruptions may go completely unnoticed, they may smoothly integrate with the user’s current task, or they may capture the user’s attention and cause a temporary or durable focus switch. Several researchers have concentrated on the effects that different notification modalities may have on the user. Robertson and his colleagues (Robertson *et al.*, 2004) analyze two types of interruptions in a debugging environment: immediate-style (i.e. interruptions that require immediate attention from the user), and negotiated-style (i.e. interruptions that the user can attend to at a chosen time). They conclude that negotiated-style interruptions are less disruptive and promote learning. McCrickard and his colleagues (McCrickard, Catrambone, Chewar, & Stasko, 2003; McCrickard & Chewar, 2003) propose to measure the effects of visual notification with respect to four parameters: (1) users' interruption caused by the reallocation of attention from a primary task to a notification, (2) users' reaction to a specific secondary information cue while performing a primary task, and (3) users' comprehension of information presented in secondary displays over a period of time, and (4) user satisfaction. They provide recommendations indicating, for example, that small in-place animation can be defined as best suited to the goals of minimal attention reallocation (low interruption), immediate response (high reaction) and small knowledge gain (low comprehension). Bartram, Ware, and Calvert (2003, p. 515) propose the use of moticons (icons with motions) as an effective, distraction-minimizing visual technique for information-rich displays. Finally, Arroyo and Selker (2003) study the effects of using different modalities for interruption in ambient displays, concentrating on the effects of heat and light channels.

2.2 Support for task switching

Current virtual (as well as physical) environments are characterized by an increasing number of resources (e.g. tools, information, communication channels) that cause users to switch between tasks very frequently. Two problems often encountered in situations of heavy cognitive load and multitasking are related to the correct continuation of planned activities, and to the evaluation of the relative priorities of concurrent tasks. These problems have also been studied in relation to *prospective memory* failures. Unlike from *retrospective memory*, which allows us to remember facts of the past (e.g. people's names, the lesson studied yesterday), prospective memory allows us to remember planned activities in the future (e.g. go to a meeting, complete writing a paper, turning off the stove in 30 minutes) (Meacham & Leiman, 1982). Retrospective memory is closely related to intentionality with some authors seeing prospective memories as the manner in which intentions are stored into memory (Marsh, Hicks, & Bryan, 1999; Sellen, Louie, Harris, & Wilkins, 1996). While prospective memory is essential to the normal progress of our daily activity, prospective memory failures may account for up to 70% of memory failures in everyday life (Kvavilashvili, Messer, & Ebdon, 2001). They have also been shown to significantly hinder performance in work and learning environments (Czerwinski & Horvitz, 2002) and to intervene differently depending on the age of the subjects (Kvavilashvili, Messer, & Ebdon, 2001).

Prospective memory doesn't simply require recalling a particular piece of information, it also requires remembering it at the *correct time*; such *correct time* may consist in an absolute time (e.g. going to a meeting at 2pm) or in the occurrence of an event (e.g. turning off the stove when the water boils). This has prompted the distinction between *event-based* and *time-based* remembering tasks (Sellen, Louie, Harris, & Wilkins, 1996).

One obvious way to support prospective memory is to supply reminder services. If a task has begun and subsequently interrupted, resuming it doesn't only require remembering to restart the task, it also entails being able to somehow re-establish its context. This may require a significant cognitive effort on the part of the user. Attention aware digital environments may support users in situations of frequent task switching by helping them in restoring the context of resumed tasks and by aiding them in recalling tasks to which they should attend. These two services are briefly discussed below.

2.2.1 Task reminders

Simple task reminders services are already available within several applications and devices. They allow users to set alarms that display the text message entered by the user when the reminder was set up. The ideal reminder service however would support both time-based and event-based remembering, supplying the user with an environment in which task reminders can be generated when, for example, a task has been completed, a resource becomes available, a document requires editing, etc. Accurate task models, associated with attention-related user models would result in better services to support users in allocating resources to pending tasks.

Task reminders may be particularly useful in helping users to remember to resume tasks that have been interrupted (a study reports that in over 40% of the cases in which tasks are interrupted, they are not resumed (O'Conaill & Frohlich, 1995)). The representation of user tasks has however presented many difficulties, often due to our poor understanding of how people form the concept of *task* as a collection of actions directed toward the achievement of a given goal.

2.2.2 Restoring task context

When a task is interrupted and subsequently restarted, a great deal of cognitive effort is spent in restoring its contexts, i.e. reassembling all the resources needed for its completion. Experimental studies have demonstrated that simple reminders about the objective of the interrupted task may be quite useful under certain conditions (Cutrell, Czerwinski, & Horvitz, 2001), however since returned-to tasks require significantly more resources, on average, than other tasks (Czerwinski, Horvitz, & Wilhite, 2004), supporting the user in recovering the resources used in performing the interrupted task would significantly lower cognitive load.

We have found that providing context restorations for resumed tasks presents a number of conceptual challenges such as:

- Establishing which resources, among those accessed while attending to the task, are significant enough to merit restoration.
- Establishing whether, at the time the task is resumed, those resources are still relevant.

In order to address these issues, we have explored the possibility of providing the learner with a multi-screen environment that supports user-guided separation of task contexts in order to facilitate resumption after task switching (Clauzel, Roda, & Stojanov, 2006). We are currently in the process of evaluating the effectiveness and usability of this system.

2.3 Self- and community-awareness tools

The increasing demands of modern learning and working environments require that users gain much greater awareness about the manner in which they allocate attentional resources. Although performance of several tasks concurrently may be improved with practice (Marois & Ivanoff, 2005), limiting multi-tasking, when applicable, is a much more efficient strategy to improve performance. This requires that users have the tools and ability to plan their activities and to reason and make decisions about their cognitive resource allocation. Awareness services informing users about their current attention-allocation choices may support such reflection. Relevant information may include details about the (type of) resources and/or tasks the user has allocated his time to, and a description of activity fragmentation (how often has a user interrupted a task, and how long did it take for him to return to it?). This information may help users in making attention-allocation decisions; for example, one may decide to block frequent sources of interruption in order to complete a task that has been frequently interrupted. Along with awareness services targeting the individual user, important insights may be gathered through community-awareness services. Notification services supporting awareness may be established to provide learners with a list of very popular resources within their community. This type of awareness tools is not limited to resource access; it may also consider a variety of community actions such as resources repurposing, bookmarking, downloading, etc.

3 An interactivist interpretation of attention

Improving the services described in the previous section requires a much better understanding of human perception and cognition, and of attentional processes in particular. We believe that the interactivist framework, by describing such mechanisms as emerging from the user's interaction with the environment, may provide the basis for answering many of the open questions stated in the following section. This section very briefly sketches how attention may be understood within the interactivist framework.

The most prominent theories in cognitive psychology see attention as the set of processes enabling and guiding the selection of incoming perceptual information. Attention limits the external stimuli processed by our bounded cognitive system and avoids its becoming overloaded (Chun & Wolfe, 2001; Driver, 2001; Lavie & Tsal, 1994; Posner, 1982). Attention can either be controlled voluntarily, or it can be captured by some external event. The former control mechanism is referred to as endogenous, or top-down, goal-driven attention (Arvidson, 2003; Posner, 1980; Yantis, 1998). The latter is referred to as exogenous, bottom-up, or stimulus-driven, and it may involve different degrees of power, whereby certain stimuli become basically impossible to ignore (e.g. sudden luminance changes), while others are more amenable to volition. Chun and Wolfe (2001, p.279) explain that "endogenous attention is voluntary, effortful, and has a slow (sustained) time course; […] exogenous attention draws attention automatically and has a rapid, transient time course". However, exogenous and endogenous mechanisms are not independent, but interact constantly so that the endogenous mechanism currently operating (e.g. what one is looking for in a visual field) may determine one's capacity to ignore certain exogenous stimuli.

We propose that within the intertactivist framework, **attention may be modeled as the choice necessary to determine which conditions in the environment are relevant for the generation of functional presuppositions**. Whenever a functional presupposition reveals itself as dysfunctional the system may learn that such dysfunction is due to an inappropriate choice of *relevant conditions* (RC).

Given a general functional presupposition FP:

(FP) *if the environment satisfies RC then interaction INT is appropriate*

if FP turns out to be dysfunctional, we have four, not mutually exclusive, possible error conditions:

- (err1) the implication FP does not hold (it is false)
- (err2) RC is not a sufficiently refined *relevant condition*
- (err3) something else was erroneously detected instead of RC
- (err4) interaction INT was erroneously enacted (somehow failed)

Correspondingly, if FP is functional, it can be reinforced at four levels: the implication holds, RC is a relevant condition, RC was appropriately detected, INT was appropriately enacted.

To a first approximation, attention may be related to error condition err2 (or the corresponding reinforcement). This error condition may trigger *learning* processes which, by monitoring ongoing interactive processes, "introduces variation when things are not going well, and stability when they are proceeding according to plan" (Bickhard, 2000, p.7-8). It is through this learning process (possibly recursively applied) that the user may generate strategies allowing selection of relevant conditions. Such *learned strategies for attention allocation* are *motivations,* as described in (Bickhard, 2000), and they emerge from a learning process about attention.

The interactivist framework provides a structure supporting the representation of attention-related information about users that is dynamically related to their past interactions with system and environment, and that takes into account the learning processes continuously underlying attention allocation. It is thus important to pursue a detailed analysis of how the interactivist framework addresses some philosophical and operational questions related to the support of human attentional mechanisms by computer systems. Such analysis may open the way to significant improvements in the design of attention aware systems.

4 Nine questions to address

The implementation of the services described in section 2 has presented difficulties largely related to the system's representation of, and reasoning about, human attentional processes. In particular, several elements of the interaction between the user and the system require more dynamic modeling. In this section we highlight some of the most pressing questions that must be addressed in order to support human attention in modern digital environments.

4.1 Question 1 - How do humans evaluate the relevance of sources of interruption?

Section 2.1.1 introduced the idea that, in order to support interruption management, a system should be able to evaluate the relevance of sources of interruptions to the user. The relevance assigned to sources of interruption by a user varies with time and context, and it is subject to a learning process based on the user's previous interactions with similar resources. For this reason, the assessment of user focus must be based both on the knowledge of the user's current environment and the relevance he may assign to elements of that environment. For example, it has become obvious that methods based solely on gaze-tracking (Baudisch, DeCarlo, Duchowski, & Geisler, 2003) are insufficient for acquiring information about users' attention. The phenomena of inattentional blindness and inattentional amnesia (Rensink, 2000) demonstrate that the current task acts as a filter on visual attention. Experimental studies (Triesch, Ballard, Hayhoe, & Sullivan, 2003) have demonstrated that the same gaze pattern may be used in the context of different tasks to acquire different visual information. Therefore, the same gaze pattern may be associated with different attentional foci. Some attempts have been made to contextualize gaze-tracking by exploiting additional sources of information (Horvitz, Kadie, Paek, & Hovel, 2003). In abstract terms, such contextualization should aim at relating two aspects of user activity. Information about the current interaction of the user with system and environment (e.g. data from gaze-tracking, mouse-tracking, observation of the user through a camera) should be related to the user's past (inter)actions and to hypotheses on how these may guide the user activity (i.e. hypotheses on possible user goals). Detecting the user’s current attentional state therefore requires a user model that dynamically describes the user state as a function of his past interactions with the environment. Interactivist models do exactly this -- they relate the user's mental states to sequences of past interactions with the environment, and therefore appear to provide the general framework necessary to create dynamic user models that integrate different sources of information about the users' (inter)actions, allowing an attention aware system to make a better guess at what the user's attentional focus might be. Further work is however necessary to better understand the processes that allow users to evaluate, at any given time, the relevance of multiple, and possibly complex stimuli.

4.2 Question 2 – How may a system move from individual interactions to bundles of intended interactions?

In section 2 we saw that an important and recurring element of the user description, as applied to all attention aware services, is the *task*. Task are used to describe the user's past, present, and future activities which, in turn, have an important influence on attention, and as such, on the choice of Relevant Conditions. But representing tasks within attention aware systems has been problematic due to the fact there are different definitions of what a task is, and different views on how a given task might be decomposed. The level of granularity at which tasks ought to be defined depends on the tasks themselves, the user, and the type of attentional support that one may wish to provide. One of the few studies on how people define *tasks* (Czerwinski, Horvitz, & Wilhite, 2004) reports that descriptions varied greatly, that people "tended to use generic terms" (ibid p. 177) for their description, and that the granularity at which tasks were defined also varied to a great extent. Given this variability, the predefined task structures used in most digital environments to date appear ill-suited to their purposes, and an emergent data definition should be sought.

Within the interactivist framework **tasks can be seen as bundles of interactions**. This implies that the user is capable, through a learning process, of acquiring and employing a new type of Functional Presupposition FP':

(FP') *if the environment satisfies RC then the bundle of interactions TASK is appropriate*

When working under this type of functional presupposition the error condition *(err1)*, under which, as we recall, *the implication FP does not hold (it is false)*, becomes more complex because its not holding is not related to a single interaction but to a bundle of interactions. Therefore the learning process generated by error condition (err1) may bring the system to revise not only the relation between RC and TASK but the structure of the TASK itself.

Knowing more about the learning mechanisms that allow a user to bundle interactions into tasks and to subsequently revise them, would help us understanding how and why people describe tasks at different level of granularity in different context (note that the choice of granularity seem to be very much related to people experience with the task).

4.3 Question 3 – How may exogenous attention be understood within the interactivist framework? Is there a difference between exogenous and endogenous attention within this framework?

It seems that from the point of view of the interactivist framework *all* changes in the environment may provoke an interaction, and that exogenous mechanisms are in fact no different from endogenous mechanisms. Both sorts of mechanisms can be seen as the result of the user's (adapted) response to its own internal state and a continuously changing environment. For a change in the environment to attract attention (as exogenous mechanisms are described) it must match some *relevant condition* recognized by the system, just as endogenous mechanisms require.

4.4 Question 4 - If an external stimulus interrupts a bundle of intended interactions, how much disruption does it provoke in the execution of those interactions?

Section 2.1.2 introduced the need to assess the level of disruption provoked by an external stimulus (interruption) at various times during task execution. Relatively recent research in visual attention has shed some light on the relation between the cognitive state of the user and the effects of interruptions/disruptions, especially with respect to the subject's ability to at least register a disruptive event (notably, see the research on change blindness (Rensink, 2000; Simons & Rensink, 2005)). However, the question of the level of disruption produced by interruptions at specific times during task execution requires further exploration. In particular, the findings of Bailey and his colleagues (Bailey, Adamczyk, Chang, & Chilson, 2006; Bailey & Konstan, 2006), by which interruptions are less disruptive when presented at the completion of coarse events, seem to imply that the smaller a bundle of interactions is, the more we are driven to complete it. This may be due to the fact that, through learning processes, users come to see certain bundles of interactions as if they were atomic interactions. A better understanding of what causes disruption, and when and why it causes it, would allow for better tuning of interruption times.

4.5 Question 5 - What role does the external stimulus modality play in the disruption of the execution of interaction bundles?

Section 2.1.3 briefly introduced various types of interruption modalities and the possible effects they have on current user activity. Most reports in the literature concern findings related to the visual modality (e.g. (Hillstrom & Chai, 2006)), and in a smaller measure to the auditory modality. Little research has so far been devoted to the possible interactions between these two, to the other modalities (but see for example (Welch & Warren, 1986)), and to their potentially different impacts on the current task. Current interactivist theories don't seem to distinguish between changes in the environment revealed through different modalities (e.g. sound, visual, tactile), however we do seem to respond to these modalities and to their possible combinations in different manners. It is possible that such different responses develop through the learning processes. Very little is known about these processes both in terms of possibly different reactions to the perception of stimuli in different modalities, and in terms of our choice of modalities for interaction with the environment.

4.6 Question 6 - How do limited human cognitive abilities impact on the selection of future interactions?

In section 2.2 we saw that limitations in memory capacity and processing ability may hinder people's ability to perform their normal activities, especially in situations of high cognitive load. Most of us have experienced the need to create "prospective memory extensions" in the form of agendas or to-do lists. It seems that when we come to the limit of our cognitive abilities we adopt "emergency strategies" that drastically reduce the number of choices we have to consider by adopting what we often call "habits". A problem often encountered when trying to support users in more efficiently allocating their cognitive resources, is that it is very difficult to modify acquired habits. For example, if a user is accustomed to regularly checking the email notification icon, he may look for that icon even when it is not needed, or even if another icon indicates that the user is working off-line. A further example is the fact that if a user has learned to ignore certain links on a web page as not useful, he is likely to keep ignoring them even when they become task-relevant. Kruschke (2001, 2003) explains these phenomena in terms of *learned attention* (highlighting) and *learned inattention* (conditioned blocking). He suggests that "both attentional shifting and associative learning are driven by the rational goal of rapid error reduction" (Kruschke, 2003, p.171) so that, in order to reduce errors with respect to already learned cues, we learn to attend to certain cues and to ignore others.

We have seen that, as understood within in the interactivist approach, the user may generate strategies allowing selection of relevant conditions, and that these *learned strategies about attention allocation* are *motivations.* Some such motivations can be identified with phenomena previously called *chronic concerns* (Moskowitz, 2002); others, generated with the aim of rapid error reduction, can be identified with the phenomena described by Krusche.

Better knowledge of the processes that guide habit creation in interaction would enable better task-reminder services and services helping the user avoiding extra load imposed by conditioned blocking and highlighting.

4.7 Question 7 - How are task and resources bundled in meaningful associations?

Our discussion of question two hinted that one of the learning processes emerging from interaction may enable a system to bundle interactions in the collections we have called tasks. In section 2.2.2 we discussed the fact that people allocate resources to tasks. In a sense, resources allocated to tasks may be *relevant conditions* for the execution of those tasks. But these conditions appear to have different epistemological status from the relevant conditions that make a given interaction appropriate. A man, for example, may recognize a possible interaction involving *throwing stone and eating*, indicated by a particular visual scan of a bird. In this case the recognition of the bird represents the relevant condition. The question is whether the availability of the stone should be considered as a further relevant condition, keeping in mind that some resources may be *necessary* for the interaction to be appropriate (as the stone in the previous example), while others may be merely *desirable* (e.g. a sling). In any case, we need to better understand the process by which resources may be bundled in collections (possibly in the form of collections of relevant conditions) that are then associated with interactions, interaction bundles, or tasks.

4.8 Question 8 - What is the impact of emotion on attentional focus?

Users may sometimes perceive attention aware services as intruding on their activities (Rudman & Zajicek, 2006) and therefore generating positive or negative emotions. Research results in behavior-processing (see (Compton, 2003) for a review) highlight the fact that fine tuning of the interface with respect to attentional processes would need to take into consideration factors such as the users' emotions, or moods. Gasper & Clore (2002), for example, report two experiments supporting the hypothesis that "affective cues may be experienced as task-relevant information, which then influences global versus local attention" (p. 34). The relevance of emotional processes in designing systems capable of interacting effectively with humans, or of simulating human behaviors, is indeed recognized by a growing community of researchers (Breazeal, 2002; Norman, Ortony, & Russell, 2003; Picard, 1997; Trappl, Petta, & Payr, 2003).

Bickhard proposes that emotions are interactions with internal dynamic uncertainty (Bickhard, 2000). As such they must influence the selection of relevant conditions (attention). Further investigation in this direction may provide very useful insights on the types of interactions an attention aware system should have with the user in order to produce intervention that are acceptable and reduce cognitive load.

4.9 Question 9 – How can a user learn from another's experiences, or from a reflection on his own experiences?

While users experience some interactions with their environment directly, in section 2.3 we saw that awareness systems may provide what might be called "mediated" and "meta-" experience. Mediated experience occurs when one is informed about others' interactive experiences, as when an adult tells a child that "the fire may burn your fingers," or when one is made aware that many others within a community have read a certain document. Meta-experience refers to a reflection about one's experience, as when one realizes, or is informed about, the fact that one has worked longer than usual hours, or that one has begun many activities and finished none. Direct experience through interaction is at the basis of the interactivist theory. Mediated experiences and meta-experiences also have an impact on the future behavior of a system, though that impact may be relevantly different.

5 Conclusions

As the interactivist framework moves from minimalist models to their more comprehensive successors it will come to address many essential questions related to human perception, cognition, and learning. This paper proposes a set of questions concerning human attention that have emerged in the study and design of digital systems capable of supporting attentional processes, and considers potentially fruitful lines of interactivist response.

6 References

Adamczyk, P. D., & Bailey, B. P. (2004). *If not now, when? The effects of interruption at different moments within task execution.* Proceedings Human Factors in Computing Systems: CHI'04, New York, 271-278.

Arroyo, E., & Selker, T. (2003). *Self-adaptive multimodal-interruption interfaces.* Proceedings International Conference on Intelligent User Interfaces (IUI'03), 6-11.

Arvidson, P. S. (2003). A lexicon of attention: From cognitive science to phenomenology. *Phenomenology and the Cognitive Sciences, 2*(2), 99-132.

Atgentive. (2005-2007). *Attentive agents for collaborative learning - eu sponsored project, fp6 - ist-4-027529-stp*.

Bailey, B. P., Adamczyk, P., Chang, T. Y., & Chilson, N. A. (2006). A framework for specifying and monitoring user tasks. *Computers in Human Behavior, 22*(4), *709-732*.

Bailey, B. P., & Konstan, J. A. (2006). On the need for attention aware systems: Measuring the effects of interruption on task - performance, error rate, and affective state. *Computers in Human Behavior, 22*(4), *685-708*.

Bailey, B. P., Konstan, J. A., & Carlis, J. V. (2001). *The effects of interruptions on task performance, annoyance, and anxiety in the user interface.* Proceedings INTERACT '01, 593-601.

Bartram, L., Ware, C., & Calvert, T. (2003). Moticons: Detection, distraction and task. *International Journal of Human-Computer Studies, 58*(5), 515-545.

Baudisch, P., DeCarlo, Duchowski, A., & Geisler, B. (2003). Focusing on the essential: Considering attention in display design. *Communications of the ACM, 46*(3), 60-66.

Bickhard, M. H. (2000). Motivation and emotion: An interactive process model. In R. D. Ellis & N. Newton (Eds.), *The caldron of consciousness* (pp. 161-178): J. Benjamins.

Breazeal, C. (2002). *Designing sociable robots*. Cambridge, MA: MIT Press.

Burmistrov, I. (2005). Interruptions in human computer interaction. Retrieved 4.1.2005, 2005, from http://interruptions.net/

Chen, D., & Vertegaal, R. (2004). Using mental load for managing interruptions in physiologically attentive user interfaces. In *Extended abstracts of the 2004 conference on human factors and computing systems* (pp. 1513-1516). Vienna, Austria: ACM Press.

Chun, M. M., & Wolfe, J. (2001). Visual attention. In E. B. Goldstein (Ed.), *Blackwell's handbook of perception* (pp. 272-310). Oxford, UK: Blackwell.

Clauzel, D., Roda, C., & Stojanov, G. (2006). *Tracking task context to support resumption (position paper for 2006 workshop on computer assisted recording, pre-processing, and analysis of user interaction data).* Proceedings HCI 2006 - Engage, London, UK.

Compton, R. J. (2003). The interface between emotion and attention: A review of evidence from psychology and neuroscience. *Behavioral and Cognitive Neuroscience Reviews, 2*(2), 115-129.

Cutrell, E., Czerwinski, M., & Horvitz, E. (2001, July 2001). *Notification, disruption, and memory: Effects of messaging interruptions on memory and performance.* Proceedings Interact 2001 IFIP Conference on Human-Computer Interaction, Tokyo, Japan.

Czerwinski, M., Cutrell, E., & Horvitz, E. (2000, 5-8 September 2000). *Instant messaging: Effects of relevance and time.* Proceedings HCI 2000 - 14th British HCI group Annual Conference, University of Sunderland, 71–76.

Czerwinski, M., & Horvitz, E. (2002). *Memory of daily computing events.* Proceedings HCI 2002 People and Computers XVI, 230-245.

Czerwinski, M., Horvitz, E., & Wilhite, S. (2004). *A diary study of task switching and interruptions.* Proceedings Proceedings of the SIGCHI conference on Human factors in computing systems, Vienna, Austria, 175-182.

Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology, 92*, 53-78.

Franke, J. L., Daniels, J. J., & McFarlane, D. C. (2002). *Recovering context after interruption.* Proceedings 24th Annual Meeting of the Cognitive Science Society (CogSci 2002), 310-315.

Gasper, K., & Clore, G. L. (2002). Attending to the big picture: Mood and global versus local processing of visual information. *Psychological Science, 13*(1), 34-40.

Gillie, T., & Broadbent, D. E. (1989). What makes interruptions disruptive? A study of length, similarity and complexity. *Psychological Research Policy, 50*, 243-250.

Hillstrom, A. P., & Chai, Y.-C. (2006). Factors that guide or disrupt attentive visual processing. *Computers in Human Behavior, 22*(4), 648-656.

Horvitz, E., Kadie, C., Paek, T., & Hovel, D. (2003). Models of attention in computing and communication: From principles to applications. *Communications of the ACM, 46*(3), 52-59.

Hudson, S. E., Fogarty, J., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. C., & Yang, J. (2003, April 05 - 10, 2003). *Predicting human interruptibility with sensors: A wizard of oz feasibility study.* Proceedings CHI 2003, Ft. Lauderdale, Florida, USA, 257 - 264.
Kaptelinin, V., & Czerwinski, M. (Eds.). (2007). *Beyond the desktop metaphor - designing integrated digital work environments*: MIT Press.
Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology, 45*(6), 812-863.
Kruschke, J. K. (2003). Attention in learning. *Current Directions Psychological Science, 12*, 171-175.
Kunz, P. (Ed.). (October 2001). *Ercim news, special issue on ambient intelligence* (Vol. 47): European Research Consortium for Informatics and Mathematics.
Kvavilashvili, L., Messer, D. J., & Ebdon, P. (2001). Prospective memory in children: The effects of age and task interruption. *Developmental Psychology, 37*(3), 418-430.
Laukkanen, J., Roda, C., & Molenaar, I. (2007). *Modelling tasks: A requirements analysis based on attention support services.* Proceedings Workshop on Contextualized Attention Metadata: personalized access to digital resources CAMA 2007 at the ACM IEEE Joint Conference on Digital Libraries, Vancouver, British Columbia, Canada.
Lavie, N., & Tsal, Y. (1994). Perceptual load as a major determinant of the locus of selection in visual attention. *Perception & Psychophysics, 56*(2), 183–197.
Mark, G., Gonzalez, V. M., & Harris, J. (2005). *No task left behind? Examining the nature of fragmented work.* Proceedings Proceedings of the SIGCHI conference on Human factors in computing systems, Portland, Oregon, USA, 321-330.
Marois, R., & Ivanoff, J. (2005). Capacity limits of information processing in the brain. *Trends in Cognitive Sciences, 9*(6), 296.
Marsh, R. L., Hicks, J. L., & Bryan, E. S. (1999). The activation of un-related and canceled intentions. *Memory & Cognition, 27*(2), 320-327.
McCrickard, D. S., Catrambone, R., Chewar, C. M., & Stasko, J. T. (2003). Establishing tradeoffs that leverage attention for utility: Empirically evaluating information display in notification systems. *International Journal of Human-Computer Studies, 58*(5), 547-582.
McCrickard, D. S., & Chewar, C. M. (2003). Attuning notification design to user goals and attention costs. *Communications of the ACM, 46*(3), 67-72.
McCrickard, D. S., Czerwinski, M., & Bartram, L. (Eds.). (2003). *Notification user interfaces special issue of the international journal of human-computer studies* (Vol. 58(5)): Elsevier.
McFarlane, D. C., & Latorella, K. A. (2002). The scope and importance of human interruption in human-computer interaction design. *Human-Computer Interaction, 17*(1), 1-62.
Meacham, J. A., & Leiman, B. (1982). Remembering to perform future actions. In U. Neisser (Ed.), *Memory observed: Remembering in natural contexts* (pp. 327-336). San Francisco: Freeman.
Moskowitz, G. B. (2002). Preconscious effects of temporary goals on attention. *Journal of Experimental Social Psychology, 38*(4), 397-404.
Nagata, S. F. (2003). *Multitasking and interruptions during mobile web tasks.* Proceedings 47th Annual Meeting of the Human Factors and Ergonomics Society, 1341-1345.
Norman, D., Ortony, A., & Russell, D. (2003). Affect and machine design: Lessons for the development of autonomous machines. *IBM system Journal, 42*(1), 38-44.
O'Conaill, B., & Frohlich, D. (1995, 7-11 May 1995). *Timespace in the workplace: Dealing with interruptions.* Proceedings CHI '95 Conference Companion, Denver, Colorado, United States, 262-263.
Picard, R. W. (1997). *Affective computing*. Cambridge, MA: MIT Press.
Posner, M. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32*, 3-25.
Posner, M. (1982). Cumulative development of attention theory. *American Psychologist, 37*, 168-179.
Rensink, R. A. (2000). Seeing, sensing, and scrutinizing. *Vision Research, 40*(10-12), 1469-1487.
Robertson, T. J., Prabhakararao, S., Burnett, M., Cook, C., Ruthruff, J. R., Beckwith, L., & Phalgune, A. (2004). *Impact of interruption style on end-user debugging.* Proceedings Proceedings of the SIGCHI conference on Human factors in computing systems, Vienna, Austria, 287-294.
Roda, C. (Ed.). (2006). *Atgentive (ist-4-027529-stp) deliverable d1.3 - atgentive conceptual framework and application scenarios*.
Roda, C., & Nabeth, T. (2006). *The atgentive project: Attentive agents for collaborative learners.* Proceedings First European Conference on Technology Enhanced Learning EC-TEL'06 (Poster section) - Innovative Approaches for Learning and Knowledge Sharing, Crete, Greece, 685 - 690.

Roda, C., Stojanov, G., & Clauzel, D. (2006). *Mind-prosthesis metaphor for design of human-computer interfaces that support better attention management.* Proceedings AAAI 2006 Fall Symposia on "Interaction and Emergent Phenomena in Societies of Agents", Arlington, Virginia, 52-59.
Roda, C., & Thomas, J. (2006a). Attention aware systems: Theories, applications, and research agenda. *Computers in Human Behavior, 22*(4), 557–587.
Roda, C., & Thomas, J. (Eds.). (2006b). *Attention aware systems. Special issue of the journal computers in human beaviour* (Vol. 22(4)): Elsevier.
Rudman, P., & Zajicek, M. (2006). *Autonomous agent as helper – helpful or annoying?* Proceedings IAT 2006 - IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Hong Kong, Forthcoming.
Sellen, A. J., Louie, G., Harris, J. E., & Wilkins, A. J. (1996). What brings intentions to mind? An in situ study of prospective memory. *Memory & Cognition, 5*(4), 483-507.
Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences, 9*(1), 16-20.
Speier, C., Vessey, I., & Valacich, J. S. (2003). The effects of interruptions, task complexity, and information presentation on computer-supported decision-making performance. *Decision Sciences, 34*(4), 771-797.
Trappl, R., Petta, P., & Payr, S. (2003). *Emotions in humans and artifacts*. Cambridge, MA: MIT Press.
Triesch, J., Ballard, D. H., Hayhoe, M. M., & Sullivan, B. T. (2003). What you see is what you need. *Journal of Vision, 3*, 86–94.
Vertegaal, R. (Ed.). (2003). *Attentive user interfaces. Special issue of communications of the acm* (Vol. 46(3)): ACM.
Weiser, M., & Brown, J. S. (1996). The coming of age of calm technology. Retrieved May 2006, from http://www.ubiq.com/hypertext/weiser/acmfuture2endnote.htm
Welch, R. B., & Warren, D. H. (1986). Intersensory interactions. In K. R. Boff, L. Kaufman & J. P. Thomas (Eds.), *Handbook of perception and human performance*. New York: Wiley.
Yantis, S. (1998). Control of visual attention. In H. Pashler (Ed.), *Attention* (pp. 223-256). London, UK: University College London Press.
Zijlstra, F. R. H., Roe, R. A., Leonova, A. B., & Krediet, I. (1999). Temporal factors in mental work: Effects of interrupted activities. *Journal of Occupational and Organizational Psychology, 72*, 163-185.