



Supporting Attention in Learning environments: Attention Support Services, and Information Management

Claudia Roda¹, Thierry Nabeth²

¹ Computer Science Department, American University of Paris 147, rue de Grenelle 75007
Paris France– croda@aup.fr, <http://ac.aup.fr/roda>

² INSEAD CALT, Boulevard de Constance, 77305 Fontainebleau Cedex, France –
thierry.nabeth@insead.edu

Abstract. Learners and knowledge workers are increasingly facing environments where frequent interruptions, multi-tasking, information overload, and insufficient community awareness are the norm rather than the exception. It has been demonstrated that this situation hinders learning in several manners. This paper, after introducing two approaches aimed at supporting attentional processes, analyses the services needed to support learners in environments presenting the above characteristics. It also discusses the conceptual and technical problems related to the collection, modelling, protection, and distribution of attention-related information.

Keywords: Attention Aware Systems, Learning Environments

1. Introduction

The advent of networked information technology has radically changed the way we value and access information, and communicate. *Information* which used to be a scarce and difficult to access resource, is now readily available, *human attention* instead has become the new “valuable currency” [15]. Taking into account this new condition when designing digital learning environments would provide learners not only access to information and people, but also the means to better manage their attention, reducing extraneous cognitive load and therefore facilitating learning [38].

In previous papers [33, 35], we have discussed findings in cognitive psychology, information science, and other disciplines that may supply an initial ground for the development of *attention aware* learning environments and we have identified the four stages necessary to support attentional processes: (1) Detecting current user’s attentional state, (2) Determining possible alternative foci, (3) Evaluating cost/benefits of possible attentional shifts, (4) Establishing modalities for interventions. Continuing in this research, in this paper we first introduce two approaches that have emerged for attention support in recent years: user-centred and resource-centred (section 2). We then analyse, in section 3, the services that should be provided to learners in order to help them coping with several attention-challenging characteristics of modern learning environments (e.g. interruptions, multi-tasking,

information overload). In section 4 we discuss the challenges related to collecting, enriching, modelling, and distributing attention-relevant information.

2. Attention support: User-Centred Versus Resource-Centred

For the purpose of attention management it is essential to model the *interactions* between people and the environment (physical environment, applications, resources, and people). Modelling such interactions requires integrating information that historically has been collected in *user* models, and *resource* models. Attention-related information may include elements such as *how, when, to which end, for how long* a learner has interacted with an application, resource, or his/her environment in general.

Two different approaches have been so far adopted to capture these interactions and provide automatic support to attentional processes: user-centred and resource-centred.

The user-centred approach is the one we have implemented in the Atgentive project [4, 34]. We model attention mainly according to a user perspective. We track and describe user activity in terms of *events* taking place in the environment (e.g. the user starting some process in an application, the user looking at the computer, the user accepting a system suggestion); based on these events a model of the user is created and system's decision are based on this user model. Within the Atgentive system we generate suggestions about possible courses of action, and provide supporting information for the current user activity (e.g. re-starting an interrupted task; interrupting the current task to perform another more urgent, or more relevant one; pursuing the current task in a different manner).

Resource-centred approaches [28, 40] model the interactions of the learner with the environment in terms of the different resources that are accessed during a learning activity. The history of resources' usage (e.g. creation, access, search, repurposing, tagging) is collected as attention-related metadata associated to individual resources; such metadata is then used to guide learners to the most appropriate resources.

Whilst the user-centred approach aims at offering a wider range of services, it has the disadvantage of requiring detailed knowledge of the user activity, which often can be gained only at the price of loss in generality. The resource-centred approach whilst aiming mainly at optimizing resource access, has the advantage of being exploitable in all the applications requiring some form of user access to resources.

In general, the information collected with a user-centred approach will also include the information that resource-centred approaches aim at collecting. For example, the fact that a user *U* starts reading a document *D*, in a user-centred approach may be represented as:

```
Event(start, task(read, U, D, ...), ...)
```

The same information, in a resource-centred approach may be represented by the following metadata associated to resource *D*.

```
<user=U, action=READ, ...>
```

Attention-relevant information about the user activity such as the user being inactive for a certain length of time can be represented in an user-centred approach, whilst such information has no meaning in a resource-centred approach (i.e. user inactivity is a property of the interaction between the user and the whole observed environment, not a property of the interaction between the user and a specific resource). Also the user-centred approach naturally represents user tasks that may "bundle" resources in meaningful associations to users' activities. For example, the fact that, user U has started writing Report1 and in order to do this he needs a certain number of resources R1, ..., Rn can be represented as:

```
Event(start, task(write, U, Report1, resources(R1, ..., Rn))
```

On the other hand, attention-relevant relationships amongst resources are better represented in a resource-centred approach. This information can provide insights on how resources are used by individuals and by the community.

These approaches are complementary in two senses. First, both information about user activity, and resource usage is necessary to supply meaningful and dynamic attention-support services (see discussion in section 3). Second, the largest set of information collected with a user-centred approach can be used to enrich resource models (e.g. by associating task related information to resources one can provide the task context in which the resource is used).

3. Services for the Support of Attention in Learning

In this section we introduce four types of services that may be provided by attention aware digital learning environments (see [31] for a complete description of the services considered within the Atgentive project) and we discuss how integrating user-centred and resource-centred approaches would greatly benefit those services' accuracy and flexibility. The services presented here illustrate various aspects of attention support in modern, community-based, and information-rich learning environments. In such environments the learners' activity is characterized by frequent interruption and multi-tasking, requiring that learners explicitly evaluate their attention allocation strategies. Interruption management is probably the service that has been most discussed in the literature and it aims at minimizing the disruption caused by the inevitable interruptions of modern learning and working environments. Support to task switching and task reminders services endeavour to help learners in managing multiple, often interleaved, tasks. Services for resources searching, ranking, and tracking aim at helping users orienteering within information-rich environments whilst minimizing information overload. Self and community awareness services provide the support necessary for the learner to assess and reflect upon his own and the community's attention allocation strategies.

3.1 Interruptions Management

Interruption management relates to the asynchronous presentation of information. For example the learner may have launched a search agent that responds asynchronously to a query, or he/she may receive an email, or a new entry may be added in a Blog that the learner is tracking. Communicating this information to the user normally results in interrupting his/her current activity. Although interruptions may bring to one's attention information possibly useful for the primary (current) task, or even, in the case of simple primary tasks, facilitate task performance [36]; it has been widely reported that interruptions increase the load on attention and memory [19], may generate stress [7, 42], and compromise the performance of the primary task [18, 26, 27, 36] especially when the user is working on handheld devices in mobile environments [27]. Therefore, in order to minimize disruption whilst ensuring that relevant 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*.

Assessing the Relevance of newly available information. Research on interruption management is quite extensive (see for example [2]), however the strategies for relevance evaluation proposed so far model data in a static manner that does not allow, for example, the generation of new categories of information to take a role in this evaluation. Although interesting results have been obtained using Bayesian models to evaluate relevance over a static structure [20], the components intervening in such evaluation should be dynamically learned by the system as it observes the user's interactions within a certain context. In [32] we propose that 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 interaction with similar resources. This suggests that by integrating resource-centred and user-centred approaches, both events describing the users' activity, and metadata associated to resources can be seen as *traces* [39] of the users' learning about relevance.

Timing of interruptions. The solutions proposed so far for the selection of interruption timing 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. Bailey and his colleagues [5, 6] 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. Alternative task decompositions have also been proposed to select interruption timing, e.g. planning, execution, and evaluation [13].

Sensory-input based timing relies on sensors' input about the user activity to detect best times for interruption. On the basis of the observation that human beings can very efficiently, and in presence of a very small number of cues, evaluate other's interruptibility, Hudson et al. [22] propose that interruptibility evaluation is attainable from simple sensors and that speech detectors are the most promising sensors. Chen and Vertegaal [10] 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. From these, user's states of interruptibility can be derived.

The integration of task-knowledge-based and sensory-input-based approaches would allow the system to rely on both types of cue for the selection of the most appropriate interruption time. Consequently, in the Atgentive project we combine knowledge of a detailed task structure [24] with simple sensory-input to evaluate the strength of breakpoints for possible interruptions.

Further integration of this user-centred approach with a resource-centred one, would enable the evaluation of the level of interruptibility necessary for the learner to cope with the complexity of the resource being proposed.

Interruptions and Collaboration. Most of the work on the evaluation of cost/benefits of interruptions has been done taking the point of view of the user being interrupted, only a few studies take into account also the cost/benefit to the interrupter, and the joint costs/benefits [21, 29]. In order to support group costs/benefits evaluation for interruptions, information about both resources use and user activity (respectively collected by resource and user centred approaches) are necessary. For example the state of development of a resource, its relevance to the community, and the deadline of the associated task, may contribute to such analysis.

3.2 Support to Task Switching

Current virtual (as well as physical) learning environments are characterized by an increasing number of resources (e.g. tools, information, communication channels) that cause learners to switch between tasks very frequently. Attention aware digital learning 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 that they should attend to. These two services are briefly discussed below.

Restoring Task Context. Learners frequently use resources as a bundle rather than individually, this is because in order to complete a task they often must access, create, and edit several different resources. When a task is interrupted and subsequently restarted, a large amount of cognitive effort is spent in restoring its contexts, e.g. 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 [12], however since returned-to tasks require significantly more documents, on average, than other tasks [14] supporting the user in recovering the resources used in performing the task would significantly lower cognitive load. In order to provide such service, user-centred approach and resource-centred approach must be combined so that the digital learning environment recognises which resources a user associates to a task at a given time of its execution.

We found that providing context restorations for resumed tasks presents a number of conceptual and technical challenges. First of all, it is necessary to establish which resources, amongst the ones originally used by the learner when attending the task, are significant enough to be restored (e.g. a learner may have accessed several web pages but probably only a subset of those pages are relevant at restoration time).

Second, it is necessary to establish the level of accuracy at which resources should be restored (e.g. should a text document that the learner was writing/reading when the task was interrupted be restored at the point where the activity was interrupted or at the beginning?) Third, given that resources are accessed through a number of different applications, how can an attention aware learning environment communicate with all other applications in order to know which resources were used, so to restore them in the appropriate manner? In order to address some of these issues, we have explored the possibility of providing the learner with a multi-screen environment that supports user-guided separation of tasks context in order to facilitate resumption after task switching [11]. We are currently in the process of evaluating the effectiveness and usability of this system.

Task Reminders. Experimental studies report that prospective memory failures (failure to remember tasks that need to be performed in the future) may account for up to 70% of memory failures [23] and that 40% of interrupted tasks are not resumed [29]. Services that help learners overcoming these problems may include simple task reminders services such as those associated to many electronic calendars. These services allow users to set alarms that display a text message entered by the user at the time when the reminder was set up. Slightly more complex reminder services are provided by shared calendars (in this case the reminder is associated to an event that has been scheduled amongst a group of people), and by calendars embedded into applications such as Microsoft Office Notification where reminders may be associated to resources, and user-created tasks. The ideal reminder service would provide the user with an environment where task reminders may be associated to user-tasks, group-tasks, as well as various types of resources. By integrating user-centred and resource-centred approaches task reminder could be automatically generated when, for example, a task has been completed, a resource becomes available, a document should be edited, etc. Further, attention-related information would allow inferring task urgency and priority making it possible to deal with the problem of task-reminders overload, which occurs when the number of reminders exceeds the capacity of the learner to attend the reminded tasks. Accurate tasks models associated with attention-related user models would result in services that better help learners in allocating resources to pending tasks. For example, current task reminder systems do not allow to distinguish (neither automatically, nor manually) between tasks that must be completed by a given date (e.g. a required assignment, an article that needs to go to print), tasks that should be completed by a given date but may be delayed if necessary, and tasks that "expire" after a certain date (e.g. go to a meeting). Furthermore, current systems only support stand-alone tasks, however tasks dependencies/sequences, and resources availability are obviously essential elements for intelligent task reminder services. For example, current systems don't allow expressing the fact that a task represents a bottleneck for other personal or community tasks, it is not possible to visualise the consequences of not completing a certain task within a certain date, and reminders are issued at pre-set times even if the conditions for the execution of a task are not met (e.g. prerequisite tasks have not been completed, or resources are not available). Intelligent task reminder services, implementing the above requirements, would lower the load on prospective memory allowing the learner to concentrate on the task currently selected.

3.3 Orienteering in Information Rich Environments

Attentional processes allow us to select, amongst incoming stimuli, those that should be further processed. These processes are obviously put under more strain as the input grows. At a low level, support to input selection in learning environments may be given by visualization tools that allow faster selection amongst several possible inputs (e.g. graphic display may enable much faster selection than textual ones). At a higher level, filtering mechanisms, such as those guiding search engines, and ranking tools, may also help the user orienteering amongst large amounts of available resources. Visualisation services are not discussed in this paper that instead briefly examines searching, ranking, and tracking of resources.

Resources Searching and Ranking. In order to help learners to find resources relevant to their needs one must have sufficient knowledge of the user, his current activity, and the available resources. For example, assume that two students in your introductory programming course formulate the same request: they want to know more about graphical user interfaces. The first student is doing very well in the course and is well on his way in the development of the assigned Java project. The second student is struggling through the course and still having problems with the analysis stage of the project. You are likely to address the first student directly to the API of some graphic library, and the second one perhaps to the tutorial of the same library.

This type of adaptive search response can be achieved only on the basis of a matching of knowledge about the students (e.g. their abilities, their current activity) and knowledge about the available resources (e.g. the original objective for their creation, the audience they are directed to, the manner how they may be, or have been, used). Similarly, a system may display adaptive search abilities only on the basis of a detailed learner model, which can be built on the basis of the information collected by user-centred approaches, and a detailed resource model, which can be collected by resource-centred approaches. These latter approaches are in fact aimed at collecting the use-history of resources. Such history provides very important insights on the context in which resources have been previously created, accessed, and modified. Erik Duval, for example, proposes "*Objects that have been used in many contexts [...] that are relevant to a specific learner, should have a higher [rank] for that learner. [...] Suppose that we track (as we can!) the correlation between the objects that learners work with and their performance on a post-test that assesses whether they have actually mastered a specific law of thermodynamics. Would that correlation not give a good indication of '[resource] quality'?*" [17].

The matching process between learner models and resource models requires dynamically extracting from the two models relevant information. For example, one may derive from the learner model the fact that a learner has already gained a certain level of knowledge (e.g. he is a good programmer) that would allow him to make efficient use of a complex resource (e.g. the API).

Whilst these types of matching have been experimented with, they are usually based on fixed categories (e.g. fast learner versus slow learner). This presents two disadvantages: First, it isn't possible to statically define all categories that may become necessary in a wide range of learning domains (e.g. good programmer versus bad programmer). Second, current solutions often require that some form of "artificial

tagging" be associated to resources and learners (e.g. the resource must be defined as being suitable for beginners, or the student must be tagged as a fast-learner). A much more effective behaviour could be obtained if instead categories were dynamically inferred from the observation of the learners, and of the resource uses in various contexts by several learners.

These types of matching may be done using data mining techniques or, in the case of smaller data sets, using heuristic based artificial intelligence techniques (see section 4.2).

Tracking Resources. Once a resource has been created, tracking services may be provided. For example, a learner may want to know when certain members of the community, notably the instructor (perhaps with a grade), have reacted to the resource. An instructor may want to know whether students have read and reacted to certain learning objects. A knowledge worker may need to know when all the members of a given group have replied to an email. Resource tracking services may provide essential input to task sequencing (e.g. one may need to wait for a set of replies to an email in order to be able to schedule a meeting). They may also provide the resource creator with knowledge about the usefulness of the resource (e.g. many accesses to the resource may indicate that the resource is useful) and therefore may guide users in the decision of whether pursuing a certain goal (e.g. keeping the resource up to date). Some of these tracking services have been implemented before (e.g. Technorati's *watchlist*) but they are not integrated in a more general environment providing attention support services and therefore they require manual user input, and the information they provide cannot automatically be reused by other services. Resource-centred approaches provide a natural information structure for resource-tracking services which can then be exploited using user-centred approaches to support task continuation and task scheduling services (these two services, not described in this paper, are part of Atgentive's conceptual framework [31] and aim at supporting learners in selecting appropriate sequences of actions when pursuing learning goals).

3.4 Self and Community Awareness Tools

The increasing solicitations of modern learning and working environments require that students and knowledge workers gain a much greater awareness about the manner in which they allocate attentional resources. Although performance of several tasks concurrently may be improved with practice [25], limiting multi-tasking, when applicable, is a much more efficient strategy to improve performance. This requires that learners have the tools and ability to plan their activity and to reason and make decisions about their cognitive resources allocation. Awareness services informing learners 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, how long did it take for him to return to it?). This information may help learners 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 but it may be based on a variety of community actions such as resources repurposing, bookmarking, downloading, etc.

Awareness services may be based on resource models as well as on task models derived from user-centred approaches.

4. Managing Attention-Related Information

Managing attention-related information presents several challenges associated to its collection, modelling, protection, distribution, and development. This section discusses each one of these challenges highlighting the issues that need to be addressed in order to develop learning environments that support attention in an efficient, un-intrusive, and secure manner.

4.1 Collecting Attention-Related Information

Although any information collection methodology in digital learning environments aims at being accurate and un-intrusive these two objectives are often difficult to achieve when dealing with attention-related information. The collection of *factual* attention-related information (e.g. the learner has accessed a certain resource three times in the last few seconds), which presents challenges in its own, often is not sufficient to guide services supporting attentional processes. It is often necessary to infer more abstract attention information (e.g. the resource accessed by the learner is related to his/her current task). Such inferred information is in general uncertain, and dynamic inference models may be necessary to derive all the interesting aspects of the learner's attentional processes.

This section concentrates on the collection of factual attention-related information, the next section briefly discusses how further information may be inferred.

Factual attention-related information may be collected in at least four manners:

- *Physical observation* collects information about the physical interaction of the user with the environment (e.g. the learner has typed 100 characters in the last minute). At the physical level, attention information is detectable by observing keyboard strikes, mouse movements, eye tracking [41], speech detectors [22], etc. Not all these methodologies are equally un-intrusive: whilst the tracking of keyboard strikes and mouse movements may be completely transparent to the user, eye-tracking devices are only recently becoming less cumbersome and require simpler setting-up procedures.
- *Psycho-physiological observation* allows collecting data about the user's psycho-physiological state (e.g. the learner's heart rate is of 160). This data can then be used to evaluate the learner's involvement with the current task or resource. Psycho-physiological information is detectable through measurements such as

Heart Rate Variability – HRV, and electroencephalogram – EEG [10]. Psychophysiological measurements are problematic because the tools necessary to take such measurements may be very intrusive; further the inferences that the system may make about the user's affective and cognitive state, on the basis of such measurements, are intrinsically uncertain.

- *Application observation* collects information about how an application is used to interact with the environment (e.g. the learner has accessed resource R using application A). Application level information is detectable by observing application software such as browsers, office applications, and learning management systems (e.g. a tool observing interaction at the application level is the Attention Recorder [3] which captures browsing history in a web browser application). Capturing attention-related information at the application level may be problematic because it requires that the application be extended to make such information available. Although this information is increasingly captured in digital environment, and even starts to be exported (c.f. the work done around APML - attention profiling mark-up language) to be exploited by external applications, asynchronous export of high level events, may not be sufficient to cover the needs of all attention related services. For example, in a text editing application, information about the position within the document where the learner pauses during a composition task may be important for establishing whether the learner is interruptible or not. Capturing this information however requires that the text editing application be capable of synchronously communicating detailed information about the current user activity with a specific resource.
- *User direct input* allows learners to provide the system with information about their states, needs, and desires (e.g. the learner has declared that he is working at assignment 1). While the above three methods of information collection are automatic and may be completely transparent to the learner, user direct input has the disadvantage of being intrinsically intrusive, and it may, in itself, be a cause of distraction from the primary task. In general however, just observing the user may not be sufficient to gain all possible types of information necessary to supply the services described in section 3. Enabling the user to augment the system observations may provide a much richer set of information. For example the learners can easily indicate their motivations and goals, whilst uncertain and complex inferences are generally necessary for the system to automatically infer such goals and motivations.

Information captured using non-intrusive automatic observations may be represented within an appropriate model and used to guide user direct input of manual annotations (e.g. resource tagging or rating) and augmentation (e.g. definition of new tasks structures). Using automatically collected information to guide manual collection has the advantages of: (1) simplifying the learner's task (for example by providing lists of already entered resource tags, and standard structures that the learner augments rather than creating) and (2) enabling consistency checks aimed at reducing error rates.

Information on task-resource association is a good example of attention related information that may significantly gain from this mixed approach; physical level and application level interactions can be captured and used to generate hypothesis on task-

resources associations, these hypothesis can be later verified with the user. Another type of data that historically has been collected manually is resource rating. Resource rating however often fails because people do not see the advantage of spending time in entering ratings. Manual ratings also have the disadvantage of easy voluntary or involuntary misuse. Resource ratings may be instead automatically derived by observation of learners' actions (e.g. bookmarks, downloads) similarly to classic social navigation techniques [9, 16, 39]. The automatically derived ratings may then be used to weight manually entered ratings.

4.2 Inferring and Enriching Attention-Related Information

In order to gain enough operable attention-related information, several types of factual information may have to be integrated. For example, methods based solely on gaze-tracking [8] are not sufficient for acquiring information about users' attention. The phenomena of inattentional blindness and inattentional amnesia [30] demonstrate that the current task acts as a filter on visual attention. Experimental studies [37] show that the same gaze pattern may be associated to different attentional foci. Attempts to address this issues include Horvitz and his colleagues' system [20] that associates gaze tracking data with information about application level activity to obtain more precise description of user foci.

In general, on the basis of factual data it is possible to derive further information about the learners' interaction with the system, their states and motivations at the time of the interaction, and about the state of the environment. Two techniques are best suited for these types of derivations: data mining, and rule based inferential reasoning. The former mechanism is necessary whenever the inference is based on a very large amount of data. The latter is applicable only when inferences are based on a small set of data. The main advantage of rule-based inference mechanisms is that rules can easily be created and modified by learners or instructors to suit their needs. On the other hand, as the number of data/facts required for the inference increases, or the set of rules increases, it is possible to assist to complexity explosion producing results that are difficult to control (e.g. rules may have conflicting effects).

Repeated inferences on factual data allow the system to enrich the attention information model with mechanisms such as:

- *Automatic tagging* (e.g. if a learner has stored a document under a folder with a certain name, the document can probably be tagged with that name).
- *Automatic resources associations* (e.g. if a learner has been copying content from a window to another, or flicking a lot between two windows, then the resources contained in the two windows are probably related; and they are probably related to the learner's current task).
- *Automatic community detection* (e.g. if two learners have directly shared – e.g. by email – many resources, they are probably close collaborators; if two learners often access the same resources, they probably share interests)
- *Default access rights definition* (e.g. if learner L1 is a close collaborator of learner L2, then L1 has visibility of all of L2's attention information and vice-versa). Rules of this type allow leveraging organizational charts and social network in order to minimise user effort in access right definition.

- *User model enrichment* (e.g. if a learner repeatedly accesses a resource, then he is probably interested in the subjects associated to it; if a learner creates several resources related to a certain subject, then he is probably an expert in – i.e. has acquired a good level of knowledge about - that subject).

4.3 Modelling Attention-Related Information

Modelling attention related information requires structures that can be dynamically augmented to consistently describe the evolution in time of the activity, needs, and interests of several learners, within several applications. In designing these models a number of constraints must be taken into account, three of the most pressing ones are briefly discussed below.

Temporal Dimension of Attention-Related Information. There is an intrinsic temporal dimension in attention-related information (e.g. interactions happen at a given time, they have a certain duration) that provides important information about attention. For example: (1) more frequently or recently accessed resources may be more relevant than those accessed less frequently or recently; (2) returning to a recently interrupted task may require less system support than resuming a task that was interrupted a long time ago.

Temporal information may also guide the integration of physical, application, and psycho-physiological observations. For example, if application P reports activity A on resource R during the time interval $t-t'$, and psycho-physiologic measurements report high levels of arousal during the same interval, one may derive that activity P on resource R (or resource R itself) is particularly important for the user.

A model of attention-related information must therefore be capable of representing time and temporal relations.

Uncertainty Management. Frequently, data mining and rule-based inferences produce results that are uncertain, this requires that sharable models of attention-related information represent uncertainty. Whilst in some situations it is possible to resolve uncertainty (e.g. by consulting the user), in others situations the uncertainty will need to be dealt-on by the systems providing services using these models.

Assessing and Modelling Tasks. One of the most relevant attention-related information about user's interaction with the environment is the task/goal of the user at the time of the interaction. Representing tasks has been however problematic due to the fact that people may have different definitions of what a task is, and how it could be decomposed. The level of granularity at which tasks should be defined depends on the tasks themselves, the user, and the type of attentional support that one may want to provide. One of the few studies on how knowledge workers may define *tasks* [14] reports great variations in how people define tasks, 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, predefined structures seem ill adapted for task description and an emergent data definition should be sought.

Whilst it is possible, under certain circumstances, to automatically capture information about users' tasks and goals, in general a much more accurate description of this type of data is obtainable if the user inputs this information manually. Users may be motivated to enter such information by the fact that they will receive high value services such as task reminders and context restoration described in section 3.

4.4 Controlling Access to Attention-Related Information, and privacy issues

Attention-related information may disclose significant details about people's activity. As a consequence, people may be willing to allow access to *their* attentional information only in a controlled manner. A significant challenge is the definition of tools enabling users to (possibly automatically) identify who will be allowed access to which attention-related information. In between the two easy choices "everyone can see all my data", and "none has access to my data" there is a possibly very large number of alternatives that, if not appropriately supported, may require a significant effort on the side of the user for their specification (e.g. user1 can access information x but not information y, user2 can access information y but not information t and z, etc.). Default access rights may be defined by using simple inferences as briefly hinted in section 4.2. The importance of protecting attention information is quite well recognised [1], and identity management systems (such as OpenID) could provide some additional insights on how to protect learner's attentional data.

5. Conclusions

Although in very recent years the importance of automatic support for attentional processes in learning environments has been increasingly recognised, a great number of both conceptual and technical challenges remain unresolved. Unfortunately, it also appears that partial solutions to attention support provide little benefit and they risk being disruptive, with the result of increasing, rather than decreasing the cognitive load on the learner. While evidence continues being collected on the negative effects of information overload, interruptions, and frequent task switching in learning environments, the knowledge provided by research in neuropsychology and cognitive psychology about attentional processes remains very hard to operationalise within the context of digital learning environments. It is still unclear exactly what information about the learners' activity is relevant for attention support and how it should be collected. Two approaches have so far been proposed to aid learners in their attention allocation: user-centred and resource-centred. However, in this paper we argue that only an integration of the two approaches may lead to the effective implementation of the complete range of services required to deal with information overload, interruptions, and frequent task switching. The wide range of learners activities, applications used in learning environments, resources and learners types, requires models that whilst capable of representing time-related and uncertain knowledge, are open and dynamic so that a shared understanding of attention-related information can be achieved, in time, across users and applications.

Acknowledgments. The work described in this paper was partially sponsored by the EC under the FP6 framework project Atgentive IST-4-027529-STP. We would like to acknowledge the contribution of all project partners.

References

1. Attention Trust, www.attentiontrust.org, accessed 16.1.2007
2. Interruptions in Human Computer Interaction, <http://interruptions.net/>, accessed 4.1.2005
3. Attention Recorder, <http://attentiontrust.org/users/atx>, accessed 13.4.2007
4. Atgentive, ist-4-027529-stp. 2005-2007.
5. Bailey, B.P., P. Adamczyk, T.Y. Chang, and N.A. Chilson, A Framework for Specifying and Monitoring User Tasks. *Computers in Human Behavior*, 2006. 22(4): p. 709-732.
6. Bailey, B.P. and J.A. Konstan, On the Need for Attention Aware Systems: Measuring the Effects of Interruption on Task - Performance, Error Rate, and Affective State. *Computers in Human Behavior*, 2006. 22(4): p. 685-708.
7. Bailey, B.P., J.A. Konstan, and J.V. Carlis. The effects of interruptions on task performance, annoyance, and anxiety in the user interface. in *INTERACT '01*. 2001.
8. Baudisch, P., DeCarlo, A. Duchowski, and B. Geisler, Focusing on the Essential: Considering Attention in Display Design. *CACM*, 2003. 46(3): p. 60-66.
9. Brusilovsky, P. Adaptive Educational Systems on the World-Wide-Web: A Review of Available Technologies. in Workshop "WWW-Based Tutoring", 4th International Conference on Intelligent Tutoring Systems (ITS'98). 1998. San Antonio, TX.
10. Chen, D. and R. Vertegaal, Using mental load for managing interruptions in physiologically attentive user interfaces, in Extended abstracts of the 2004 conference on Human factors and computing systems. 2004, ACM Press: Vienna, Austria. p. 1513-1516.
11. Clauzel, D., C. Roda, and G. Stojanov. Tracking task context to support resumption (Position paper for 2006 Workshop on Computer Assisted Recording, Pre-Processing, and Analysis of User Interaction Data). in *HCI 2006 - Engage*. 2006. London, UK.
12. Cutrell, E., M. Czerwinski, and E. Horvitz. Notification, disruption, and memory: Effects of messaging interruptions on memory and performance. in *Interact 2001 IFIP Conference on Human-Computer Interaction*. 2001. Tokyo, Japan.
13. Czerwinski, M., E. Cutrell, and E. Horvitz. Instant messaging: Effects of relevance and time. 14th British HCI group Annual Conference. 2000. Univ. of Sunderland, UK.
14. Czerwinski, M., E. Horvitz, and S. Wilhite. A diary study of task switching and interruptions. in *Proceedings of the SIGCHI conference on Human factors in computing systems*. 2004. Vienna, Austria: ACM Press.
15. Davenport, T.H. and J. Beck, *The Attention Economy*. 2001: Harvard Business School.
16. Dieberger, A., P. Dourish, K. Hřšk, P. Resnick, and A. Wexelblat, Social Navigation: Techniques for Building More Usable Systems. *interactions*, 2000. 7(6): p. 36-45.
17. Duval, E., LearnRank: the Real Quality Measure for Learning Materials (Tematic Dossier 06 December 2005), *Insight - Oservatory for New Technologies and Education*.
18. Franke, J.L., J.J. Daniels, and D.C. McFarlane. Recovering context after interruption. in *24th Annual Meeting of the Cognitive Science Society (CogSci 2002)*.
19. Gillie, T. and D.E. Broadbent, What makes interruptions disruptive? A study of length, similarity and complexity. *Psychological Research Policy*, 1989. 50: p. 243-250.
20. Horvitz, E., C. Kadie, T. Paek, and D. Hovel, Models of attention in computing and communication: from principles to applications. *CACM*, 2003. 46(3): p. 52-59.

21. Hudson, J.M., J. Christensen, W.A. Kellogg, and T. Erickson. I'd be overwhelmed, but it's just one more thing to do: availability and interruption in research management. in SIGCHI conference on Human factors in computing systems. 2002. Minneapolis, USA: ACM press.
22. Hudson, S.E., J. Fogarty, C.G. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J.C. Lee, and J. Yang. Predicting Human Interruptibility with Sensors: A Wizard of Oz Feasibility Study. in CHI 2003. 2003. Ft. Lauderdale, Florida, USA: ACM Press.
23. Kvavilashvili, L., D.J. Messer, and P. Ebdon. Prospective memory in children: The effects of age and task interruption. *Developmental Psychology*, 2001. 37(3): p. 418-430.
24. Laukkanen, J., C. Roda, and I. Molenaar. Modelling tasks: a requirements analysis based on attention support services. in Workshop on Contextualized Attention Metadata: CAMA 2007 at the ACM IEEE Joint Conference on Digital Libraries. Vancouver, Canada.
25. Marois, R. and J. Ivanoff. Capacity limits of information processing in the brain. *Trends in Cognitive Sciences*, 2005. 9(6): p. 296.
26. McFarlane, D.C. and K.A. Latorella. The Scope and Importance of Human Interruption in Human-Computer Interaction Design. *Human-Computer Interaction*, 2002. 17(1): p. 1-62.
27. Nagata, S.F. Multitasking and interruptions during mobile web tasks. in 47th Annual Meeting of the Human Factors and Ergonomics Society. 2003.
28. Najjar, J., M. Wolpers, and E. Duval. Attention Metadata: Collection and Management. in WWW2006 workshop on Logging Traces of Web Activity: The Mechanics of Data Collection. 2006. Edinburgh, Scotland.
29. O'Conaill, B. and D. Frohlich. Timespace in the Workplace: Dealing with Interruptions. in CHI '95 Conference Companion. 1995. Denver, Colorado, United States: ACM press.
30. Rensink, R.A., Seeing, sensing, and scrutinizing. *Vision Research*, 2000. 40(10-12): p. 1469-1487.
31. Roda, C., ed. Atgentive (IST-4-027529-STP) Deliverable D1.3 - Atgentive conceptual framework and application scenarios. 2006.
32. Roda, C. Supporting attention with dynamic user models (Extended abstract). in Interactivist Summer Institute 2007. Paris.
33. Roda, C. and T. Nabeth. The role of attention in the design of Learning Management Systems. in IADIS International Conference CELDA (Cognition and Exploratory Learning in Digital Age). 2005. Lisbon, Portugal: IADIS Press.
34. Roda, C. and T. Nabeth. The AtGentive project: Attentive Agents for Collaborative Learners. in First European Conference on Technology Enhanced Learning EC-TEL'06 (Poster section). 2006. Crete, Greece: Springer.
35. Roda, C. and J. Thomas. Attention Aware Systems: Theories, Applications, and Research Agenda. *Computers in Human Behavior*, 2006. 22(4): p. 557-587.
36. Speier, C., I. Vessey, and J.S. Valacich. The effects of interruptions, task complexity, and information presentation on computer-supported decision-making performance. *Decision Sciences*, 2003. 34(4): p. 771-797.
37. Triesch, J., D.H. Ballard, M.M. Hayhoe, and B.T. Sullivan, What you see is what you need. *Journal of Vision*, 2003. 3: p. 86-94.
38. van Merriënboer, J.J.G., P.A. Kirschner, and L. Kester, Taking the load of a learners' mind: Instructional design for complex learning. *Educational Psychologist*, 2003. 38(1): p. 5-13.
39. Wexelblat, A. and P. Mayes. Footprints: History rich Web browsing. in Conference on Computer-Assisted Information Retrieval. 1997.
40. Wolpers, M., G. Martin, J. Najjar, and E. Duval. Attention Metadata in Knowledge and Learning Management. 6th Intern. Conf. on Knowledge Management. 2006. Graz, Austria.
41. Zhai, S., What's in the eyes for attentive input. *CACM*, 2003. 46(3): p. 34-39.
42. Zijlstra, F.R.H., R.A. Roe, A.B. Leonova, and I. Krediet, Temporal factors in mental work: Effects of interrupted activities. *Journal of Occupational and Organizational Psychology*, 1999. 72: p. 163-185.