2 Human attention and its implications for human–computer interaction

Claudia Roda

Remembering planned activities, resuming tasks previously interrupted, recalling the names of colleagues, sustaining focused performance under the pressure of interruptions, ensuring that we don't miss important information . . . these are only a few examples of critical activities whose performance is guided by attentional processes. This chapter proposes that knowledge about attentional processes can help us design systems that support users in situations such as those described above. The first part of the chapter gives an overview of some of the essential theoretical findings about human attention. The second part analyses attentional breakdowns and how those theoretical findings may be applied in order to design systems that either help avoid attentional breakdowns or assist in recovering from them.

2.1 Introduction

Current information and communication technologies concentrate on providing services to users performing focused activities. However, focused activity is no longer the norm. Users are often interrupted, they switch between the contexts of different devices and tasks, maintain awareness about the activity of distant collaborators and manage very large quantities of information. All this results in high cognitive load that may hinder users' overall achievements.

In order to address interaction in a more realistic manner, we have been working on the development of systems that are capable of supporting the processes that govern human cognitive resources allocation: attentional processes.

Attention plays an essential role in task performance and interaction. It enables us to act, reason and communicate, in physical or virtual environments that offer us stimuli exceeding, probably by several orders of magnitude, what we are actually capable of processing. Attention makes it possible for us to pursue goals without being distracted by the immense variety of available alternative stimuli and actions and undeniably mediates our interaction with the world.

Many years of research, within several fields of study, have demonstrated that attention is a surprisingly complex and multifaceted phenomenon. However, as we discover more about the processes involved in attention, we are also increasingly provided with the knowledge necessary to design systems that take into account the limitations and characteristics of such processes. This is particularly important because people interact with a growing number of devices while involved in many parallel activities. Hence the strategies and means employed for allocating and shifting attention play a major role in performance and satisfaction.

In our approach, the essential cues enabling the understanding of user activity are the user interactions with the environment. Such interactions are managed by attentional processes, which guide the allocation of cognitive and physical resources, allowing one to both perceive the environment and act upon it. Attention allocation can be used as the proxy that both reveals and guides interactions enabling us to build *attention-aware systems* (Roda and Thomas 2006). These systems recognize that attentional processes play an important role in many of the problematic situations faced by users of digital environments and aim at reducing information overload, limiting the negative effects of interruptions, increasing situation awareness (especially in the case of virtual environments) and supporting users in situations of multi-tasking (Roda and Nabeth 2007). In our work, for example, we have been able to show that attention management may effectively guide interaction in digital learning environments. The results obtained show that attention-based scaffolding improves students' results, while fostering a more proactive attitude towards the learning activity and increased motivation (Molenaar and Roda 2008 and Molenaar *et al.* in chapter 11 of this book). Similar results highlighting the positive effects of attention support have been obtained by others in situations of cooperative problem solving (Velichkovsky 1995) and in contexts where the user needs proactive assistance (Eisenhauer *et al.* 2005).

One problem that has often been encountered in designing attention-aware systems is that current knowledge about the cognitive and perceptual processes underlying attention allocation is, if seen from an HCI (human–computer interaction) point of view, very scattered. At the macro-level, many different theories, based on diverse hypotheses, describe individual aspects of attention, but no unified view of attentional phenomena exists. At the micro-level, research results about individual attentional phenomena are often analysed for very simple tasks and environments which, while allowing for sound and well-controlled experimental settings, do not reflect at all the conditions of users in real-world applications. Unfortunately this situation is not likely to change in the short term. The integration of the different aspects of attention in a single theory capable not only of describing individual phenomena but also of predicting their effects and interactions seems currently out of our reach. Perhaps easier to achieve is the scaling-up of some of the findings reported on individual phenomena so that they are a closer approximation of real-world settings in which users select their own goals, read documents composed of many words, see screens whose content depends on previous operations, etc.

The aim of this chapter is to collect the findings of psychological research that appear most relevant to the design of attention-aware systems (section 2.2) and then to show how these findings have been, or could be, used in design (section 2.3). Given the breadth of this review, it is necessarily very partial, but it will hopefully give the reader a feeling for the issues involved in designing systems that take into consideration human cognitive and perceptual limitations.

We set the scene with the classic endogenous versus exogenous perspective on attention and then explore two important areas of study: divided attention and automaticity. Understanding divided attention is essential to the design of attention-aware systems because, under this heading, we find research highlighting the constraints under which we perform multiple tasks and attend to multiple sensory input. Automaticity, on the other hand, explores what we appear to be able to do more easily, although the subsection on ‘what we may miss’ mitigates the view of our efficiency. Section 2.2 concludes with an overview of the important relationship between attention and memory and a discussion of long-term attention which is almost completely excluded from current studies in cognitive psychology and neuroscience. In section 2.3 we turn to the application of psychological theories to system design. In order to do this, we consider common situations of failure, which we name *attentional breakdowns*, and describe how attention-aware systems may help avoid, or recover from, such breakdowns. In particular we consider: prospective memory failures; retrospective memory failures; task resumption failures; disruption of primary tasks; missing important information; and habituation errors. In discussing recovery and avoidance of these breakdowns, we consider several types of systems; however, we don’t discuss here three large application domains: machine vision, robotics and virtual reality. We believe that most of the discussion in this chapter would also apply to these domains, but a treatment of their specific requirements is outside the scope of the chapter.

2.2 The many faces of attention

Attention has been extensively studied for many years. However, the answer to the question *what is attention?* is not a straightforward one. *Attention as selection* has been the most common paradigm guiding research in this field (Baddeley and Weiskrantz 1993; Driver 2001; Lavie and Tsal 1994; Parasuraman and Davis 1984; Posner 1982), although some authors stress that attention selectivity covers a variety of very different purposes and functionalities (see, for example, Allport 1993). Within the *attention as selection* paradigm attention is seen as the set of mechanisms that allows the allocation of cognitive resources, which are assumed to be limited. In the literature, attentional selection has been associated with a variety of – possibly overlapping – functions, including influence over (1) which stimuli will be processed, (2) which information will enter working memory (Awh, Vogel and Oh 2006; McNab and Klingberg 2008), (3) which stimuli will reach a level of conscious availability (Koch and Tsuchiya 2007; O'Regan and Noë 2001; Posner 1994) and (4) which internal and external actions will be performed (Hommel 2010; Hommel, Ridderinkhof and Theeuwes 2002; Norman and Shallice 1986).

With respect to visual attention, for example, Desimone and Duncan (1995: 194) summarize attentional selection as follows: 'At some point (or several points) between input and response, objects in the visual input compete for representation, analysis, or control. The competition is biased, however, towards information that is currently relevant to behaviour. Attended stimuli make demands on processing capacity, while unattended ones often do not.' With respect to action, Norman and Shallice (1986: 3) propose that 'two complementary processes operate in the selection and control of action. One is sufficient for relatively simple or well-learned acts. The other allows for conscious, attentional control to modulate the performance.'

This section discusses three aspects of attention that are particularly relevant to HCI. First, in section 2.2.1, we are concerned with the issue of how attention may be affected by the environment and by the internal state of the user (e.g., his goals, intentions, motivation) and how these effects may interact. This knowledge will provide us with a better understanding of how, by acting on the user environment, devices may direct or *protect* users' attention. Second, in section 2.2.2, we explore how attention may be divided among several targets. This aspect of attention is obviously related to multi-tasking, which is a normal condition of operation in most computing environments. The objective is to gain an understanding of how the organization and presentation of several tasks and information may affect user performance. Third, in section 2.2.3, we consider the issue of automaticity. Automatic processes are those that can take place without disturbing ongoing activity. If a device can communicate with users by activating automatic processes then the communication is very efficient and does not disturb the user. Fourth, section 2.2.4 explores the relationship between attention and memory through two constructs: working memory and prospective memory. The former has often been correlated with intelligence; it significantly impacts on the efficiency with which we can treat information and defines the limits to the amount of information we can elaborate at one time (Buehner *et al.* 2006; Conway *et al.* 2002; Engle 2002; Engle, Kane and Tuholski 1999; Engle, Tuholski *et al.* 1999). The latter controls our ability to perform planned actions; because of its high failure rate, supporting prospective memory is particularly important. Finally, section 2.2.5 briefly discusses the time span of attention over which digital support takes place.

2.2.1 Endogenous/exogenous – top-down/bottom-up processes

Attention selectivity can be considered as guided by two main mechanisms. Either attention is captured, in a ‘bottom-up’ manner, by external events – as when one notices a sudden loud noise in the silence – or it is controlled voluntarily, in a ‘top-down’ manner, by the subject – as when one follows the sequence of words in a text one is reading. The two types of control are often called respectively exogenous and endogenous to stress the fact that either external or internal (to the subject) events regulate attention allocation. This dichotomy, bottom-up versus top-down, is in many ways related to the classic dichotomy, recurrent in twentieth-century psychology, focusing on either conscious control of human behaviour, as proposed by humanist theories, or behaviour which is determined by environmental factors, as in early behaviourist theories, and unconscious choices, as proposed by Freud. Many current theories of attention assume that both aspects intervene, so that some human experiences and behaviours are automatic responses to environmental stimuli, whilst other experiences and behaviours are under the control of the subject. The top-down, bottom-up dichotomy has also been the source of a debate related to the fact that some authors see attention as a *cause*, others see it as an *effect*, and others yet as a combination of both (Fernandez-Duque and Johnson 2002; Stinson 2009). Under the causal interpretation, attention is seen as an engine capable of orienting perception and guiding cognitive processes. Such a motor is generally modelled through some ‘executive system’ which, some authors dispute, is none other than a homunculus because no clear account is given of its functioning. Effect theories of attention, instead, see attention allocation as the result of various sensory and cognitive processes. These theories, rooted in neuroscience, maintain that no executive system exists and perceptual stimuli compete in order to activate cortical areas, and attention is merely a side effect of these competitive processes. So, while cause theories associate attention with top-down processes and dispute whether attention plays a role in bottom-up processes as well (i.e., whether there can be any processing of sensorial input without attention), effect theories merely see attention as a by-product of bottom-up processes (i.e., attention plays no role in the processing of sensorial input). Whilst the main objection to cause theories is the homunculus issue, the main objection to effect theories is their alleged inability to account for situations in which very salient stimuli are not attended, or vice versa, low-saliency stimuli are.

As we return to the discussion of top-down (or endogenous) and bottom-up (or exogenous) processes, we will see that, although this chapter mainly reports on causal theories, the themes mentioned above will recur often.

An important difference between the two attentional mechanisms is that exogenous processes are assumed to be capable of processing several stimuli in parallel, while endogenous processes are considered to be sequential; consequently the former are much faster than the latter. Chun and Wolfe (2001: 279) stress the fact 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'.

The interaction between exogenous and endogenous processes has been the subject of much research and it is often studied through models based on the observations of subjects' physical and/or neurological activity. Following most theories, overall attentive behaviour cannot be determined by one or the other type of processes individually. However, from the point of view of HCI, it is important to note that exogenous processes are triggered by changes in the environment, i.e., something a device may be able to provoke, whereas endogenous processes are under the subject's internal control which a device may only be able to influence indirectly.

Following this classic differentiation between endogenous (top-down) and exogenous (bottom-up) processes, many authors have proposed more detailed models describing how these processes may work.

Bottom-up processes select stimuli on the basis of their *saliency*, where saliency is determined by how much an item stands out from its background based on *basic features* (e.g., colour, shape, etc.), luminance, level of detail or extended configurations (Rensink, chapter 3 in this volume). Other factors that appear to influence bottom-up selection may be learned – e.g., hearing one's own name in a conversation is very salient, and a famous face generates more interference than an unknown one (Lavie 2005) – or are instinctively important – such as translating and looming stimuli (Franconeri and Simons 2003) or novel signals (Fahy, Riches and Brown 1993). Note, in passing, that this strictly bottom-up definition of saliency is not shared by all authors. Bowman *et al.* in chapter 5 of this book, for example, define saliency in terms both of bottom-up and top-down processes, including factors such as relevance to long-term goals and emotional significance.

Top-down processes, instead, select stimuli on the basis of their relevance to the current task or goal. This selection may be done by enhancing the quality of the signal of stimuli that have certain task-relevant features at a given time. Top-down processes are based on information describing which characteristics of the input are relevant to the current task. Duncan and his colleagues call this information *the attentional template* (Desimone and Duncan 1995; Duncan and Humphreys 1989). It also appears that the strength of the bias associated with certain input characteristics 'depends on the difficulty of the task performed at the attended location' (Boudreau, Williford and Maunsell 2006: 2377) so that, for example, if a stimulus is more difficult to recognize, the top-down signal supporting its selection will be stronger.

An important aspect of selective attention is related to the control of action. In order to explain how action may be controlled, including the cases in which action performance may be considered automatic, Norman and Shallice propose that two different and complementary sets of processes are involved. The first set of processes controls actions that are 'relatively simple or well learned' (Norman and Shallice 1986: 3); in this case, action sequences are represented by sets of *schemas* that may be activated or inhibited by perceptual input without the need for attention. Different levels of activation enable the selection of schemas through a mechanism called *contention scheduling*. The second set of processes depends on a *supervisory attentional system* (SAS) and provides for the management of novel or complex actions for which no schema is available. The SAS intervenes by supplying extra activation or inhibition of schemas so that the appropriate sequence of actions may be selected that responds to the situation.

This model fits well with the bottom-up, top-down paradigm described earlier. Sensory-based (bottom-up) and volition-based (top-down, involving the SAS) activation processes interact to guide action.

Along with Norman and Shallice's, several other models have been proposed which aim to articulate this interaction between attention, perception, consciousness and action (e.g., Hommel, Ridderinkhof and Theeuwes 2002; LaBerge 2002).

Based on results of functional neuroimaging, Posner and his colleagues propose that three distinct functions of the attentional system should be recognized: alerting, orienting and executive control. 'Alerting is defined as achieving and maintaining a state of high sensitivity to incoming stimuli; orienting is the selection of information from sensory input; and executive attention involves mechanisms for monitoring and resolving conflict among thoughts, feelings, and responses' (Posner and Rothbart 2007: 7; see also Posner and Fan 2007; Hussain and Wood 2009).

Within this framework we can imagine that signals such as alarms and warning road signs would vary the state of alertness; the provision of spatial cues for where a target will appear would orient attention; and executive control may be activated when planning is needed, to detect errors (e.g., attention is needed for one to realize that one has chosen the wrong road), to respond appropriately to novel situations or to overcome habitual actions (e.g., typing on an English qwerty keyboard when used to a French azerty one).

The analysis proposed by Posner and his colleagues provides important insights for human–device interaction. The first of these is the existence of a general alertness state that would make a user more sensitive to incoming stimuli. Second, there is the possibility of using cue-based orienting of attention to support users in making selections without reducing available choices (see section 2.3.4 of this chapter). Third, there is the need to take into consideration the increased effort the user will have to invest in *novel situations* and in *overcoming habitual actions* (see section 2.3.6).

As a result of the activation of bottom-up and top-down processes, a selection takes place that enables only the strongest signals to influence subsequent processing. Note that this type of selection in fact happens at many levels between sensory input and higher level processing.

In certain situations bottom-up priority may be so high that a signal takes over attention even if it is irrelevant to the current task. The involuntary shift of attention to a target that is not relevant to the current task is called *attention capture* (Franconeri and Simons 2003; Yantis 2000). The issue of whether attention may be captured in a purely bottom-up manner, and what exactly are the characteristics of the stimuli that may trigger such a capture, is still a subject of research: see Gibson *et al.* 2008 for an account of the many aspects and interpretations of *attention capture*. It is clear, however, that under certain conditions, certain stimuli – e.g., sudden luminance changes or noise – cause a shift of attention in a manner that appears to be independent of the current task. Lavie and her colleagues propose a theory that aims at clarifying the different roles played by perception and cognitive control in attention capture. On the basis of a set of experiments, they argue that high perceptual load reduces distractor interference, whilst high cognitive load increases distractor interference (Lavie *et al.* 2004).

Attention capture is very important for the design of human–device interfaces because, on the one hand, devices may be able to 'protect' users from undesired attentional shifts (e.g., someone's phone ringing in a lecture theatre may distract a whole audience), but, on the other hand, devices may be able to provoke attention capture when a user's attention needs to be drawn to a particular event (e.g., calling an operator's attention to a fault in the system he is controlling): see section 2.3.4 of this chapter.

2.2.2 Divided attention

Attention may be concentrated on a single item (focused attention) or it may be divided between multiple targets (divided attention, split attention). The majority of the work on divided attention addresses one or both of two related issues: (1) multi-tasking, in particular dual-task performance; and (2) the identification of multiple sensory inputs. In both cases divided attention has been shown frequently to induce errors and delays in response. The questions addressed are: Which cognitive processes are involved in the performance of two or more tasks simultaneously or in attending multiple sensory inputs? And what are the factors intervening in the performance of multiple tasks? The answers to these questions have important consequences for device design and for how information should be presented in order to facilitate learning (this latter aspect is discussed by Low, Jin and Sweller in chapter 4 of this volume).

Two main theories have tried to explain the problems we may encounter in divided attention situations: capacity theories and cross-talk theories. Another hypothesis which is relevant to divided attention is that multi-tasking involves switching from one task to another and that the switch itself may generate interference.

2.2.2.1 Capacity theories Capacity theories argue that a limited pool of cognitive resources is available. Some authors postulate that we have a single set of mental resources (Kahneman 1973) and, consequently, as we increase the number of targets, we necessarily reduce the resources available to attend each one of them. Other theorists argue for a *multiple* *resources theory* by which different cognitive and perceptual processes are supported by different sets of resources and therefore performance under divided attention varies depending on whether the targets require the same resources or not. Wickens (2002) identifies four types of resources (dimensions) influencing task interference: processing stages, perceptual modalities, visual channels and processing codes.

The processing stage dimension predicts that perceptual and cognitive activities share the same resources while selection and execution of responses pulls from a separate set.

The perceptual modalities dimension predicts that different perceptual modalities (visual, auditory, etc.) pull from separate resources. For example, Duncan and his colleagues (Duncan, Martens and Ward 1997) found that targets in different modalities do not generate the same level of interference as multiple sensory input presented in the same modality. Note that, although multiple task performance is obviously affected by the limit of perceptual analysis of multiple stimuli, there are situations, such as split visual attention over easily discriminated targets, in which it appears possible simultaneously to attend stimuli at non-adjacent locations (Bichot, Cave and Pashler 1999; Cave and Bichot 1999; McMains and Somers 2004).

The visual channels dimension predicts that focal vision requires a different set of resources than ambient vision.

Finally, the processing codes dimension predicts that analogue/spatial processes use a different set of resources than categorical/symbolic (e.g., linguistic) processes.

2.2.2.2 Cross-talk theories Cross-talk theories attribute the errors and delays that one may experience in divided attention situations not to the fact that there is, so to speak, not enough fuel to support multiple cognitive activities, but rather to the interference between the contents of the information being processed. These theories relate performance to the information involved in the specific tasks, so that similar tasks are more likely to interfere with each other. Several experiments show that dual-task performance improves when the two tasks are dissimilar. Navon and Miller (1987: 435) report experimental results supporting the hypothesis that reduced performance in dual-task situations may be due to interference when 'the outcome of the processing required for one task conflicts with the processing required for the other task (e.g., cross-talk)'.

2.2.2.3 Task switching Multi-tasking is also closely related to task switching (Pashler 2000). Many experiments demonstrate that if two tasks must be attended in sequence, the response to the second task is slowed down as the interval between the two tasks is reduced. This effect is termed the *psychological refractory period* (PRP) (Welford 1952). One possible explanation of the delay observed when people try to divide attention between two or more tasks is that only one active task-set (i.e., the configuration of mental resources necessary to perform the task (Anderson 1996; Monsell 2003)) can be maintained at a time. Under this hypothesis, multi-tasking amounts to frequent switches of attention between the attended tasks. The task-set is changed at each switch. The multi-tasking activity is therefore affected by the delayed response times due to the PRP. The PRP, and task-switching delays in general, have been extensively studied and several alternative explanations of this effect have been proposed (Meiran, Chorev and Sapir 2000; Pashler 1994; Pashler and Johnston 1998). Rogers and Monsell (1995) present a set of experiments indicating that both task-set updating costs and cross-talk effect intervene in task switching.

Altmann and Trafton (2002) have performed experiments on a task requiring frequent switches between goals (the Tower of Hanoi puzzle) and formulate a *goal-activation* model. The main hypotheses guiding this model are that goals have different levels of activation in memory, that decay of memory traces is not instantaneous but gradual, and that the most active goal is the one that will guide behaviour. The authors argue that three elements can be used to predict performance: first, the interference between goals due to decay time for old goals in memory; second, the time needed to encode the new goals; and third, the cues available for retrieving pending goals. We will see that these three predictive constraints play an important role in the design of attention-aware systems, in particular with respect to prospective memory failures and disruption of primary task.

2.2.2.4 Diffusion of attention Recent research has reported an opposite effect of divided attention that, although frequently experienced, has rarely been studied. Exploring the *attentional blink*, an effect by which subjects fail to identify the second of two visual targets presented in close succession, Olivers and Nieuwenhuis (2005: 265) found that this effect 'is significantly ameliorated when observers are concurrently engaged in distracting mental activity, such as free-associating on a task-irrelevant theme or listening to music'. In order to explain these results they formulate the hypothesis that the task-irrelevant mental activity generates a *diffusion of attention* which could be attributed to a higher state of arousal, a positive affective state or the multi-tasking situation itself.

Another study reports that complex choices (e.g., deciding which car or apartment to buy) may actually benefit from the lack of attention, and subjects may achieve more satisfactory results if, during the decision-making time, their attention is engaged in an unrelated demanding task. These results were explained by the fact that conscious thought, which can be very precise, is also limited by the boundary of what we can attend to at any given time. Unconscious thought, instead, can process and summarize very large amounts of information (Dijksterhuis *et al.* 2006).

2.2.3 Automaticity

The discussion in the previous section has highlighted the existence of two types of processes: those that can, in a sense, be considered *automatic* and those that require a closer control on the side of the subject. Automaticity is pervasive in human behaviour and extends to the automatic effect of perception on action, automatic goal pursuit and a continual automatic evaluation of one's experience (Bargh and Chartrand 1999). In this chapter, however, we will only concentrate on two aspects of automaticity that are particularly relevant to HCI: the lower effort required to perform automatic processes as compared to non-automatic ones, and the high effort required to override automatic reactions. Although the discussion so far has given an indication that bottom-up processes are automatic whilst top-down ones are not, a clearer definition of what automaticity is and which processes actually correspond to this definition would be helpful.

In the literature automaticity has been defined in many different ways and factors of very diverse nature have sometimes been considered. In fact automaticity can be defined along at least three different sets of parameters: (1) the behaviour induced; (2) the neuronal mechanisms involved; and (3) the cognitive mechanisms underlying the processes. So, for example, on a behavioural basis we can say that an automatic process will induce a fast response to a stimulus, on a neuronal basis we can say that amplified activity takes place in a certain area of the brain, and on a cognitive mechanism basis we can say that the process does not require the intervention of an executive attentional system. With respect to our objectives, the distinction between the three sets of parameters is important because behavioural and neuronal parameters enable us to give a measurable definition of the occurrence of automatic processes. This means that on the basis of behavioural or neurophysiologic observations of the user we will be able to predict the likelihood that certain environmental conditions will trigger automatic processes. In particular we will concentrate on behavioural parameters, which can give us a sense of how users may respond to certain types of interaction.

Historically, several different behavioural parameters have been used to define automatic processes. First, automaticity is normally associated with fast *response times*. Response time measures the time interval between the presentation of a stimulus and the response of the subject (e.g., pressing a button in response to seeing a certain object on the screen). Second, automaticity is normally associated with obligatory execution, i.e., the subject may not be able to avoid executing the process. Third, automatic processes are assumed to have no interaction with other concurrent processes, i.e., in situations of divided attention, the performance of other processes is not affected by the automatic process. Fourth, it has been argued that automaticity is normally associated with high transferability so that the performance level of automatic processes remains constant across different types of tasks. Fifth, automaticity is normally associated with no awareness, i.e., the subject will not be able to report that the process is taking place. Note also the relation between choice and awareness: no awareness requires that the process be obligatory. Sixth, automaticity is normally associated with no sensitivity to distractors so that the presence of multiple stimuli will not affect the level of performance of the automatic process.

Most authors consider only the first two of these parameters and automatic processes are defined as obligatory processes resulting in fast response times.

Although these parameters have frequently been defined as taking discrete values (e.g., processes are parallel and fast or serial and slow, interaction takes place or not), there is increasing evidence that they may take continuous values, so that a process may generate a continuum of response times or may interact with other processes at different levels under different conditions.

Behavioural and neuronal measures are often used to deduce underlying cognitive mechanisms. However, there is no widespread agreement on what combinations of such measures imply which cognitive mechanisms. For example, in order to assess whether certain subjects' responses to stimuli are purely due to bottom-up cognitive mechanisms (i.e., purely controlled by external stimuli irrespective of the subjects' attentional state), many experiments rely on *response time* or *stimulus exposure duration*. Following the experimental technique employed by Treisman and Gelade in defining *feature integration theory* (Treisman and Gelade 1980), many authors consider that if the response time is relatively short and constant, unrelated to the number of distractors, then the process is bottom-up (and pre-attentive in particular) because subjects are obviously not performing a serial search through the items but a parallel one. Similarly, for stimulus exposure duration, pre-attentive processing is assumed to take place when subjects to whom a stimulus is shown for a short and fixed exposure duration (about 200 ms) accurately report on the presence of the target stimulus, regardless of the number of distractors. However, these types of definitions have caused confusion between pre-attentive processes and, what we will call here, learned-response processes which, along other dimensions of automaticity, behave significantly differently from pre-attentive processes (see section 2.2.3.2).

2.2.3.1 On what we perceive fast Early selection theories (Broadbent 1958) and modified early selection theories (Treisman 1960), which are briefly discussed in the introduction to this book, stipulate that essential information about sensory input is extracted by one type of automatic processes, pre-attentive processes, and is then processed by attentive processes. Pre-attentive processes are defined as bottom-up processes dealing with simple information about the input signals and, importantly, they are very fast because input is processed in parallel (Treisman 1985; Wolfe 2001; see also Rensink, chapter 3 of this volume).

Response time in pre-attentive processing is fast and not significantly affected by the size of the display, and it can take place when focused attention is prevented (e.g., by the simultaneous performance of an attention-demanding task, or by extremely brief exposure to the stimulus). According to Treisman and Gelade's feature integration theory (Treisman and Gelade 1980), pre-attentive processes are the bottom-up processes that detect *basic features* of the visual input, such as colour, orientation and size. It is still a matter of research which *basic features*, and under which conditions, are systematically detected by pre-attentive processes: see Wolfe 2001 for an overview.

Several authors (see, for example, Logan 1992: 317; Wright 1998: 111) define pre-attentive processes as being obligatory, stimulus-driven, parallel, independent of attention and preceding attentional selection: the output of pre-attentive processes is assumed to be the input for attentive ones. However, a classic question in the attention literature is whether, and if so which, bottom-up processes really act independently of top-down control, i.e., can take place in the absence of attention.

In the visual modality Treisman (Treisman 1985; Treisman and Gelade 1980) has proposed that individual features are processed in parallel and, at a later stage, attention intervenes to *integrate* these features into objects. Consequently, searches for targets defined by individual features are parallel and not affected by variations in the number of distractors, whilst searches along multiple features require attention and are therefore sequential. For example, the red spot and the diagonal bar in figures 2.1a–d (see plate) can be found pre-attentively, both targets *pop out* and their retrieval time is not significantly different whether they are surrounded by few distractors (figures 2.1a and c) or many (figures 2.1b and 2.1d). On the other hand, attention needs to be applied to find the blue diagonal bar in figure 2.1e and the response time will be much higher in the case of a larger number of distractors as shown in figure 2.1f. The red spot and the diagonal bar of figures 2.1a–d are said to be *visually salient* because they stand out from their background of homogenous blue spots in figures 2.1a and 2.1b, or horizontal bars in figures 2.1c and 2.1d.

Similarly to the visual modality, research in the sound modality has demonstrated that parallel processing supports the recognition of auditory features such as frequency, intensity and duration of acoustic stimuli (see, for example, Takegata *et al.* 2005; Winkler *et al.* 2005).

Other authors, however, note that visual saliency is not an absolute property of a stimulus but it describes how a certain element stands out with respect to its background (Itti 2005). For example, the saliency of the diagonal bar in figures 2.1c and 2.1d is significantly reduced amongst non-homogenous distractors in figure 2.1g, and amongst distractors that are very similar to the target as shown in figure 2.1h. On the basis of these observations, Duncan and Humphreys (1989: 433) suggest that in visual searches 'difficulty increases with increased similarity of targets to nontargets and decreased similarity between nontargets producing a continuum of search efficiency'. This is a departure from the classic dichotomy (feature/parallel versus conjunctions/serial) governing theories about visual searches that stipulate that searches for individual features proceed in parallel, whilst searches for conjunction of features take place serially. Automaticity is no longer simply associated with searches for individual features. The quality of the distractors, and not only whether the search is for single/multiple features, comes into play in deciding whether the search can be performed fast, obligatorily and without interacting with other processes. A related question is that of how different features contribute to overall perceptual saliency. Several computational models of bottom-up attention have represented different featural contributions and control mechanisms (see Itti 2005 for a review), providing a better understanding of feature interaction. These models are based on *saliency maps* representing stimulus saliency in every point of a two-dimensional space: see, for example, Itti, Koch and Niebur 1998. Recently, models integrating both bottom-up and top-down attentional control have been developed (e.g., Navalpakkam and Itti 2006; Schill, Zetzsche and Hois 2009). These models not only contribute to the field of machine vision but also enable the testing of hypotheses about the functioning of human attentional mechanisms.

Another line of research departing from feature integration theory has aimed at establishing what other types of information, beside pre-attentive features, can be extracted automatically from images. Evidence from neuroscientific observations, for example, supports the hypothesis that fairly complex facts, such as determining whether an image contains the picture of an animal, can be extracted very quickly (within 150 ms) upon presentation of a stimulus (Thorpe, Fize and Marlot 1996) and, in general, the gist of a scene can be determined very rapidly (Rensink, chapter 3 of this volume). Other work suggests that in some cases discrete objects may be recognized with the same efficiency as individual features (see, for example, review in Scholl 2001).

Although the application of these laboratory results to interface design is not always straightforward, they imply that certain information can be made available to the user by presenting it in a manner that triggers automatic processes, thereby minimizing the demands on the attentional system (as discussed in section 2.3.4, this would be particularly useful for notification systems). In fact, the discussion so far would place us in an ideal situation, allowing the design of interfaces that present information in a manner that does not disturb the user and where interaction could be mostly automatic, requiring little or no attention from the user.[1] Unfortunately, as our experience tells us, this is not always possible. In the following sections we will see that attention may be necessary to detect stimuli even if they are very salient, and that the performance of concurrent tasks may be negatively affected even by well-known interaction patterns.

2.2.3.2 On what we can learn to do rapidly In the classic perception literature automatic processes have often been equated to pre-attentive processes and defined as processes that do not require attention. More recent accounts (Logan 1992; Treisman, Vieira and Hayes 1992) distinguish between pre-attentive processes and what we here call *learned-response* processes. Learned-response processes are associated with learning. The idea is that performing a well-rehearsed action (e.g., recognizing a well-known object, eating with fork and knife, washing your hands) will require less attentional effort than performing a new or less known one. Logan believes that this difference in cognitive demand may be explained by a theory in which ‘novice (nonautomatic) performance is based on a general algorithm for solving the problems the task presents, whereas

---

[1] Note that interfaces controlled mostly on automatic feedback are not unusual and include, for example, the dashboard and control system used to drive a car.

automatic performance is based on single-step, direct-access retrieval of past solutions from memory . . . Automatic processing has the properties of well-practiced memory retrieval. It is fast and effortless' (Logan 1988; 1992: 321). This dual nature of automaticity is also evident in the analysis of automatic self-regulation proposed by Bargh and Chartrand, as they state: 'Some of the automatic guidance systems . . . are "natural" and don't require experience to develop . . . Other forms of automatic self-regulation develop out of repeated and consistent experience; they map onto the regularities of one's experience and take tasks over from conscious choice and guidance when that choice is not really being exercised' (Bargh and Chartrand 2000: 476).

For our purposes it is important to note that there appear to be two types of processes allowing fast access to information. First, there are pre-attentive processes (such as single-feature visual searches), which are either innate or acquired at a very early age. These processes maintain performance in situations of divided attention and in the context of different tasks. Second, there are learned-response processes (such as the recognition of some danger symbols in a display), which may become automatic but only after practice. These processes bring no advantages to performance in divided-attention settings (Logan 1992) and don't transfer well to novel tasks (Treisman, Vieira and Hayes 1992).

This implies that in display design certain pop-out effects can only be achieved through training whereas others, based on single features, come *for free* as they are innate.

2.2.3.3 On what we may miss We have seen that a question motivating the research work mentioned above is that of whether automatic processes, and bottom-up processes in general, require attention, i.e., the intervention of some executive attentional system. Whilst the work presented in sections 2.2.3.1 and 2.2.3.2 implies that attention may not be involved in early processing of sensorial input, some neuroscientific experiments support the hypothesis that attention operates from the very early stages of visual processing (Awh, Vogel and Oh 2006; Hillyard, Vogel and Luck 1998), with amplified responses to attended visual stimuli beginning within 60 ms of stimulus onset. The results of some behavioural experiments may also be explained as refuting the hypothesis that certain visual processes occur *without attention* (Gibson and Peterson 2001). Several authors (Mack and Rock 1998; Simons and Chabris 1999), for example, argue that there is no conscious perception of the visual world without attention to it (but see also discussion in Driver *et al.* 2001). Mack bases her argument on the *inattentional blindness* phenomenon which 'denotes the failure to see highly visible objects we may be looking at directly when our attention is elsewhere' (Mack 2003: 180). Several experiments demonstrate that highly salient (from a sensory point of view) stimuli can be completely missed if they are not the explicit targets of a visual search; therefore it is argued that, unless attention is allocated to the target, the subject develops no conscious perception of the stimuli.

The possibility of a subject completely missing certain parts of a stimulus is accompanied by another, similar phenomenon, *change blindness*, i.e., the 'failure to see large changes that normally would be noticed easily' (Simons and Rensink 2005: 16). Changes in the visual environment are normally salient because they produce a transient motion or flicker. However, in a series of, sometimes very surprising, experiments (Simons) the authors show that even very large changes may go unnoticed if they are not attended when they occur (see also Rensink in chapter 3 of this volume).

Interestingly, the factors that may prevent change detection are not only related to the current attentional focus (a subject concentrating on a target may miss a change occurring in the environment), or sensory input (e.g., the change is hidden by an occluding object, the flicker of a display or an eye movement), but could also be cultural. Nisbett and Masuda, for example, report that East Asian subjects are more likely to detect changes in the relationships between objects in a scene, whereas Westerners are more likely to detect changes to objects' attributes (Nisbett and Masuda 2003).

Taken together, the main result of this research is that human vision does not create a *copy* or *complete representation* of the world in the mind, as has been assumed for many years. Human vision rather seems to be a more dynamic process that binds elements of the external world in models that satisfy the needs of the viewer on the basis of the current task. As O'Regan puts it, 'the outside world is . . . a kind of external memory store which can be accessed instantaneously by casting one's eyes (or one's attention) to some location' (O'Regan 1992: 461).

2.2.3.4 Automaticity in action So far we have mainly concentrated on the role of attention in the selection of perceptual input. The question of automaticity, however, naturally highlights another important aspect of attention: the role it plays in action control. In their seminal paper studying this aspect of attention, Norman and Shallice argue that the term *automatic*

> has at least four different meanings. First, it refers to the way that certain tasks can be executed without awareness of their performance (as in walking along a short

> stretch of flat, safe ground). Second, it refers to the way an action may be initiated without deliberate attention or awareness (as in beginning to drink from a glass when in conversation). Third, it is used in cases such as the orienting response, in which attention is drawn automatically to something, with no deliberate control over the direction of attention. And finally, [it refers] to situations in which a task is performed without interfering with other tasks' (Norman and Shallice 1986: 1–2).

As we have seen earlier, it may happen that the sensory-based activation is strong enough to override volition-based activation. In these cases actions that may be unrelated or even inappropriate for the performance of the current task may be initiated or completely carried out (as when one walks to a place out of habit when, in fact, one should have gone somewhere else). These situations correspond to the first two types of automaticity described by Norman and Shallice. As the authors note, the third type of automaticity – 'attention is drawn automatically to something' – is significantly different from the previous two types because, instead of guiding action without involving attention allocation, it automatically redirects attention.

We may therefore (in the first two cases of automaticity) have stimuli that provoke certain actions but don't involve any attentional shifts. The subject maintains his attention on the current task. In the third case, however, the automatic process triggers a change in the supervisory system that may provoke a lasting change of attentional focus. In terms of human–device interaction, if we assume that some device produces the stimulus, the latter case of automaticity corresponds to the generation of a very salient stimulus that attracts attention to the device itself (for example, to provide information about an emergency situation). The former two cases instead correspond to a stimulus capable of producing an automatic action that will not disturb (or, more likely, bring minimal disturbance to) the user's current activity such as when, for example, one stops or starts walking at the change of a traffic light. This type of automatic behaviour requires some learning on the side of the user but the resulting interaction is very efficient.

2.2.4 Attention and memory

Human thought and action are obviously influenced by past experience. Memory is the system that enables us to record past experience and use it in the present. Attention is strictly related to memory in two manners: on the one hand, our memories may influence attention allocation; on the other hand, attention allocation may determine which sensory input is stored in memory.

The study of human memory has a long history and memory has been defined along many different dimensions.

One classic differentiation is between explicit memory and implicit memory: see Eichenbaum 1997, Polster, Nadel and Schacter 2007 and Schacter 1987 for influential reviews. Implicit memory is not available for conscious retrieval but influences task performance; one example of implicit memory is procedural memory which enables us to perform actions, such as riding a bicycle, without explicitly remembering the individual components of the action. Explicit memory enables us to bring facts and experiences of the past to mind and then express them in some format. The memory of facts such as *Rome is the capital of Italy*, which are theoretical knowledge independent from a specific context, has been named *semantic memory*; the memory of personal experiences, which are related to specific contexts in time and space and carry some emotional value, has been named *episodic memory* (Tulving 1972).

Although the distinction between implicit and explicit memory is important and has influenced research in human–computer interaction (e.g., Oulasvirta 2004; Oulasvirta, Kärkkäinen and Laarni 2005), this chapter will mainly focus on two other characteristics of the memory system that are more immediately related to attention: the relationship between long-term memory and working memory, and the distinction between retrospective memory and prospective memory.

2.2.4.1 Working memory Memory is normally seen as fulfilling two different functions: collecting information for long-term retrieval and holding information for immediate usage. These two different functions can be exemplified by the memory of the number 313 that may be created to remember the date when Constantine issued the Edict of Milan; and the memory of the number 313 created if asked to add 213 and 100. The former type of memory is called *long-term* memory (LTM), referring to the fact that the number 313 is stored for retrieval possibly hours, days or even years after it is memorized. The latter type of memory is meant for immediate use and the number 313 (which, by the way, may be retrieved from long-term memory as the date of the Edict of Milan) is only remembered for the time necessary to perform the calculation. Memory for immediate use is characterized by fast decay (it does not last long) and very limited capacity which, in a very influential paper, Miller (1956) evaluated at 7±2 *chunks* of information (for a more recent discussion of working-memory capacity limits, see also Cowan *et al.* 2008).

In the seventies Baddeley and Hitch (1974) proposed a model of this memory for immediate use, which they called *working memory*. Since then the term working memory has been adopted widely. Baddeley and

Hitch emphasized the fact that working memory includes several components, that it combines processing and storage, and that it forms the basis for most cognitive activities. The authors proposed that working memory is a three-component system, including 'a control system of limited attentional capacity, termed the *central executive*, which is assisted by two subsidiary storage systems: the *phonological loop*, which is based on sound and language, and the *visuospatial sketchpad*' (Baddeley 2003: 830). Subsequently Baddeley (1986) proposed that the central executive system could be implemented with the model proposed by Norman and Shallice (1986) (also discussed in sections 2.2.1 and 2.2.3.4 of this chapter).

The Baddeley and Hitch model has been fundamental for the development of much research addressing how attention and working memory relate to other aspects of cognition, including language learning and processing, fluid intelligence, consciousness and many other cognitive processes (Engle 2002). The model has also been supported by much neuroscience research, notably by the findings of Goldman-Rakic (1987). However, some authors have argued for a dynamic, functional view of working memory, by which working memory is the 'active portion of LTM, coupled with mechanisms for cognitive control' (Conway, Moore and Kane 2009: 262). In this view, the content of working memory may lose its discrete characteristic and, as Anderson has proposed, because of 'the continuous nature of activation . . . membership in working-memory is a matter of degree. Less active working-memory elements are processed less rapidly, for instance, in a recognition task' (Anderson 1983: 263). Items are available for processing, not because they are stored in a special component, but because they have reached a threshold value of activation in long-term memory. In this view, control mechanisms of working memory are not achieved through a specialized system but rather due to 'coordinated recruitment, via attention, of brain systems that have evolved to accomplish sensory-, representation-, or action-related functions' (Postle 2006: 23).

Cowan's model (Cowan 1988) integrates some aspects of the classic Baddeley and Hitch model with the view of degrees of activation in long-term memory. Cowan argues that what really distinguishes short-term from long-term memory is the processes necessary to maintain activation, e.g., rehearsal for short-term memory, and semantic elaboration for long-term memory. Memory can be activated either automatically by external stimuli, or through attention. These two types of activations may interact so that 'automatic activation may direct attention, and attention may in turn influence the amount of memory activation' (Cowan 1988: 172). Further, Cowan argues that the *focus of attention* is a part of the active portion of memory, and that items in the focus of attention are immediately available for processing. Cowan's model has subsequently inspired similar models that include a *focus of attention* component: see, for example, Oberauer 2002.

Complementary to activation, another aspect of memory is very important for the good functioning of cognition, that is decay. Decay reduces activation and eventually results in forgetting (or losing the memory trace). Altmann and Gray (2000) argue that, in task environments, we need to forget a task (at least partially) in order to perform another one. In particular, if one task has reached a below-threshold state before a new one becomes current, then there will be no interference between the two tasks. Literature in this area has concentrated on addressing the question of what are the factors that intervene in forgetting: see Wixted 2004 for a survey.

2.2.4.2 Prospective memory Prospective memory, closely related to intentionality (Marsh, Hicks and Bryan 1999; Sellen *et al.* 1997), is the mechanism that allows us to remember planned activities in the future (e.g., go to a meeting, complete writing a paper, turn off the oven in thirty minutes, give a message to a friend when we meet him). Whilst retrospective memory is the mechanism that allows us to remember facts of the past (e.g., people's names, the lesson studied yesterday), prospective memory requires remembering to remember, i.e., remembering something at an appropriate moment in the future. Such a moment may be represented by an actual time (e.g., going to a meeting at 2 p.m.) or by the occurrence of an event or a series of events (e.g., publishing the minutes once everyone has approved them). This has brought about the distinction between 'event-based and time-based remembering tasks' (Sellen *et al.* 1997: 484). Studies in prospective memory have mostly concentrated on event-based remembering and analyse the two aspects of 'acting when encountering the correct circumstances (prospective component) and . . . remembering the correct action to perform (retrospective component)' (Kardiasmenos *et al.* 2008: 746).

Kliegel *et al.* describe prospective memory mechanisms as organized in four phases: '(1) intention formation – the point at which a future activity is planned; (2) intention retention – the period during which the intended action is retained in memory while other ongoing activities are performed (i.e., the ongoing task); (3) intention initiation – the moment at which execution of the intention is initiated; and (4) intention execution – the actual execution of the intended action(s) according to a previously formed plan' (Kliegel, Mackinlay and Jäger 2008: 612). As in the case of other cognitive processes discussed above, the question arises of the cognitive effort required in order to complete prospective memory tasks. In particular, several authors have studied how one may go from intention retention to intention initiation. This process, often named *prospective memory retrieval*, requires matching some event in the environment to target events in prospective memory (e.g., it's 2 p.m., thirty minutes have elapsed, all approvals to the minutes have arrived) and initiating the task to be executed (e.g., go to the meeting, turn off the oven, publish the minutes). Some authors (e.g., Burgess and Shallice 1997) see prospective memory retrieval as an attention-demanding process and have proposed that the executive attentional system (such as Norman and Shallice's supervisory attentional system or Baddeley's central executive) monitor the environment for events that would match prospective memory target events. Other authors have proposed that prospective memory retrieval may be an automatic process that does not require attention allocation except perhaps in specific situations such as when the task is very important. In the latter interpretation, events in the environment may act as cues that, rather than requiring focused attention, are accepted by an automatic-associative memory and, if activated enough, may bring to awareness the associated intended action (Einstein and McDaniel 1996).

'Critical for purposes of prospective memory, this information is retrieved rapidly, obligatorily, and with few cognitive resources [and,] in contrast to the cue-focused views, the target event is not necessarily recognized as a cue' (McDaniel *et al.* 2004: 606). However, a set of experiments proposed by McDaniel and his colleagues (McDaniel *et al.* 2004) support the hypothesis that both types of processes may intervene in prospective memory retrieval. In particular, they argue that retrieval is more likely to be attained through automatic – i.e., reflexive and obligatory – processes if the target event is sufficiently associated with the intended action (e.g., write the word 'house' whenever the picture of a house appears), whilst attentive processes are more likely to intervene when the target event is not as well associated with the intended action (e.g., write the word 'house' whenever the number 56 appears). There appears to be no single rule deciding the type of process applied in retrieval; rather, the degree of monitoring versus spontaneous retrieval depends on task complexity and individual differences.

Einstein (Einstein and McDaniel 1996) found that factors intervening in the selection of automatic or attentive processes include: (1) whether the target event is focal, i.e., whether the ongoing activity encourages focal processing of the prospective memory target; (2) the importance of the prospective memory task, as reflected by the level of emphasis given to the prospective memory task when instructing the subjects;

(3) the number of target events; (4) the duration of the ongoing task; and (5) individual differences. In their experiments, spontaneous retrieval processes took place with focal prospective memory target events and when the prospective memory task was moderately emphasized. On the other hand, attentive monitoring of the target event took place 'with a nonfocal target or high-emphasis instructions or both' (Einstein and McDaniel 1996: 331). A single target event was retrieved automatically, whilst a test with six target events revealed that attentive monitoring was engaged. Consistently with previous results indicating that capacity for maintaining controlled processing is limited (e.g., Bargh and Chartrand 1999), the authors found that attentive monitoring declined over trials in the nonfocal condition. Optimal performance of the prospective memory task, with minimal costs (least disturbance to the ongoing task), were obtained in the focal target moderate-emphasis condition.

2.2.5 Long-term and short-term attention

Our perception, intentionality, social situation and aesthetic sensibility all seem to concur in determining our attentional behaviour. This behaviour may be observed and analysed both in the short term and in the long term. Short-term attentional processes, strongly related to working memory and cognitive load, reflect one's immediate concentration on an object or activity. For example, short-term attention may be deployed on a specific part of an image in order to recognize an object. Long-term attentional processes refer to processes that span over a length of time of minutes, hours, days or even months. They normally involve one or more tasks, long-term memory as well as working memory. An example of long-term attention may be the cognitive effort that one makes in writing a letter, or in completing a much longer, possibly collaborative, project. Understanding short-term attention enables one to support users in the immediate selection of, and focus on, tasks or objects. Short-term attention can be evaluated by the use of behavioural and psychophysiological measures capable, for example, of detecting the level of arousal of a user in relation to a given stimulus.

Understanding long-term attention enables one to provide individuals and groups of users with the appropriate information and guidance about their long-term allocation of cognitive resources. Long-term attention can be inferred through analysis of subjects' activities over an extended period of time. To our knowledge, long-term attention has not been investigated in psychology and neuroscience. The only research that moves in the direction of long-term attention involve those experiments requiring the performance of lengthy complex tasks under high cognitive load conditions. Examples of such tasks are complex-span tasks (Diamond 2005; Unsworth and Engle 2007) where subjects are asked to remember a list of items while performing some processing activity, e.g., calculations. Complex-span tasks are sometimes called working-memory tasks because they test the ability to maintain (or focus on) elements held in working memory. Some experiments (e.g., Colom *et al.* 2006) compare results obtained by subjects performing complex-span tasks with results obtained in simple-span tasks (or short-term memory tasks) where subjects have to remember lists of items but are not asked to perform a simultaneous processing activity. Conway and his colleagues (Conway *et al.* 2002) indicate that simple-span tasks can be performed on the basis of automatic routines such as rehearsal and chunking, whilst complex-span tasks cannot and therefore require the intervention of the supervisory attentional system.

2.3 Addressing attentional breakdowns

In the previous section we explored some of the current theories of attention which may play a significant role in the design of interactive systems. This section explores a set of situations in which breakdowns occur as a consequence of particularly demanding conditions, which in modern working environments are often due to multi-tasking and interruptions (Czerwinski, Horvitz and Wilhite 2004; Gonzalez and Mark 2004; Mark, Gonzalez and Harris 2005). Activity fragmentation stresses attentional processes, long-term memory and working memory, and has been the subject of several studies. For example, Gonzalez and Mark report that ‘In a typical day . . . people spend an average of three minutes working on any single event before switching to another event [and] somewhat more than two minutes on any use of electronic tool, application, or paper document before they switch to use another tool’ (Gonzalez and Mark 2004: 119).

In this section we will suggest how systems have been, or could be, designed in order to avoid, or recover from, attentional breakdowns.

In the human–computer interaction literature one finds references to attention with respect to very short time spans, as in the examples given in section 2.2.3, as well as to much longer time spans, as in the attention needed to drive a car (C. Ho, Tan and Spence 2005; Pêcher, Lemercier and Cellier 2009), perform collaborative activities (Nabeth and Maisonneuve in chapter 12 of this volume) or learn academic subjects (Molenaar *et al.* in chapter 11 of this volume). These situations are characterized by the fact that people are often involved in activities that require multiple and possibly interdependent or closely related tasks, they have to interact with many people and devices, and they have to find, create or manage very large amounts of information.

As we move from short to longer time spans, attentional processes are increasingly seen as coordinating information flow (rather than just selecting information input) so that appropriate higher-level perception and action can take place. Research in developmental cognitive neuroscience, for example, suggests that problems in initiating and monitoring task performance, in both adults and children, are due to deficits in the *central executive* (Diamond 2005; Wilson *et al.* 1998). Executive functions of attention are therefore necessary to plan activity and establish priority between multiple competing tasks. These observations are obviously based on a causal interpretation of attention. Effect theories, however, point in the same direction. For example, with respect to visual attention, Rensink states that attention is 'the establishment (and maintenance) of a coordinated information flow that can span several levels of processing' (Rensink 2007: 139).

This double aspect of *attention as selection* and *attention as coordination* is consistent with literature that sees attentional breakdowns as the cause of failure in situations that may appear very diverse. Forgetting to start or complete a task, being prone to interrupting one's primary task, difficulties in restarting an interrupted task or in establishing whether the conditions for the execution of a task are met, problems with task prioritization and inability to find and focus on relevant information, are all examples of phenomena that have been attributed to attentional breakdowns. This section analyses these failure situations and proposes how appropriate system design may reduce some of the burden on the user's cognitive system. We define the following set of typical attentional breakdowns: (1) prospective memory failures; (2) retrospective memory failures; (3) task resumption failures; (4) disruption of primary task; (5) missing important events and information; and (6) habituation-related failures. Each breakdown situation is analysed in detail.

2.3.1 Prospective memory failures

As discussed earlier, normal activity often requires remembering, at appropriate times, plans we have made in the recent or distant past. These memories enable the correct continuation of planned tasks when they have been interrupted, and the evaluation of relative priorities of concurrent tasks. In section 2.2.4.2 we saw that prospective memory is the mechanism that enables these types of recollection. Daily experience and laboratory experiments demonstrate that prospective memory is essential for professional performance, independent living and social relationships (Eldridge, Sellen and Bekerian 1992; Kardiasmenos *et al.* 2008; Lamming *et al.* 1994). However, it has been reported that prospective memory failures may account for up to 70 per cent of memory failures in everyday life (Kliegel and Martin 2003; Kvavilashvili, Messer and Ebdon 2001) and that such failures are likely to occur more frequently in older adults than in younger ones (Kliegel, Mackinlay and Jäger 2008; Kvavilashvili, Messer and Ebdon 2001; Zimmermann and Meier 2010). This situation is further aggravated in modern working and learning environments where a high level of multi-tasking increases the difficulty in keeping track of relative priorities between tasks.

Services to help users overcome these problems may include task reminder services such as those associated with many electronic calendars. The ideal reminder service, however, should provide the user with an environment where task reminders may be associated with user tasks and group tasks, as well as various types of resources. These reminders should also help users remember to resume tasks that have been interrupted, which are reported as not being resumed in 40 per cent of cases (O'Conaill and Frohlich 1995).

One approach to enhancing current reminder services is based on collecting information about people's attention allocation to task and resources (see Schmitz *et al.*, chapter 8 of this volume on such information collection) and then using this information for inferring tasks' urgency, relationships and priority. Reminder systems should be able to represent and manage information about tasks' dependencies/sequences and resource availability. This would enable the system to detect if a task represents a bottleneck for other personal or community tasks, to visualize the consequences of not completing a certain task within a certain date, and to issue reminders only if the conditions for the execution of a task are met (e.g., prerequisite tasks have been completed and resources are available). Intelligent task reminder services, implementing the above requirements, would lower the load on prospective memory, allowing users to concentrate on the task currently performed.

With respect to interrupted tasks, the *goal-activation model*, briefly introduced in section 2.2.2.3, predicts that prospective memory failures could be reduced 'if operators were taught to react to an alert by searching for a cue and associating it with the goal being suspended' (Altmann and Trafton 2002: 66). The authors also argue that the digital environment should provide those cues so that an association can be formed at interruption time and priming can take place at resumption time.

In section 2.2.4.2 other factors affecting the effort associated with prospective memory tasks were discussed. Some of these factors, such as the emphasis on importance of the prospective memory tasks, cannot be controlled through system design. Others, however, may be. These include high association between target and intended action, focal targets and single target. Under these conditions prospective memory tasks may be achieved with higher levels of automaticity, therefore reducing cognitive effort.

It seems likely that systems capable of providing semantically relevant reminder messages highly associated with the reminded tasks will prove effective in reducing users' cognitive effort in the phases of *intention initiation* and *intention execution*. Further, integration of the target (i.e., the reminder message) within the current task environment would make the target focal, and could result in the improved performance of prospective memory tasks. Finally, the aggregation of all reminder messages in a single system should promote automation since the user will be monitoring a single target area or object type. To our knowledge none of the above hypotheses has been verified through the evaluation of interfaces providing these types of services. Much work, however, has been done in the evaluation of notification methodologies (see section 2.3.4), of which reminders are a special case. It should be noted that in the design and evaluation of notification systems the emphasis is often placed on reducing the disturbance to the primary task whereas here we have been concerned with techniques that optimize the retrieval of a prospective memory task.

2.3.2 Retrospective memory failures

Retrospective memory failures occur when someone has difficulties remembering previously acquired information (e.g., someone's name, a lecture studied, having met someone, having visited a place). Eldridge proposes that retrospective memory problems can be classified into the following seven categories: forgetting a person's name, forgetting a word, forgetting an item in a list, forgetting a past action or event, forgetting some aspect of past actions or events, forgetting where some object (physical or electronic) was put or last seen and forgetting how to perform some action or series of actions (Eldridge, Sellen and Bekerian 1992). Several systems have been devised with the aim of reducing the negative effects of retrospective memory failures or enhancing human memory capacity (e.g., Mase, Sumi and Fels 2007). The 'Forget-Me-Not' system (Lamming *et al.* 1994; Lamming and Flynn 1994), for example, has been implemented in a PDA-like device and continuously collects data about user activity (e.g., telephone calls, documents printed, people encountered, etc.), allowing the user to search information for specific events.

For instance, one can retrieve the telephone number dialled while talking to a colleague. Because associations of this type (i.e., not remembering a telephone number but remembering that a call to that number was made in a certain context) happen frequently, services such as those offered by the 'Forget-Me-Not' system promise to be very useful to remedy retrospective memory failures. iRemember (Vemuri and Bender 2004), a 'wearable memory prosthesis', captures audio – and face-to-face conversations in particular – along with the user's location, calendar, email, commonly visited websites and weather. This data is then indexed and organized for searching and browsing so that retrieved data can act as triggers for forgotten memories.

Issues related to retrospective memory failures have concerned researchers for a long time and recur across the entire field of information management (e.g., Freeman and Gelernter 2007; Karger 2007).

2.3.3 Task resumption failure: context restore

Because multi-tasking is the condition under which we often operate, we regularly have to interrupt tasks and restart them at a later time. Mary Czerwinski and her colleagues performed a diary study analysing the effects of interruptions on task performance. They found that it is significantly more difficult to switch to tasks that require 'returning to' after an interruption, and that returned-to tasks take generally longer than more routine tasks and require 'significantly more documents, on average, than other tasks' (Czerwinski, Horvitz and Wilhite 2004: 178–9). Mechanisms for prospective memory support, such as the ones described in section 2.3.1, may be used to remind a user of the need to resume interrupted tasks. Resuming a task, however, doesn't only require remembering to restart the task but also entails being able to re-establish the context of that task (e.g., retrieving the documents necessary for the performance of the task). To this end, retrospective memory support systems such as those mentioned in the last section can be used to remind the user about the context of resumed tasks – see Czerwinski and Horvitz 2002 and Franke, Daniels and McFarlane 2002 for a review from this perspective, and Franke, Daniels and McFarlane 2002 for an example of a system implementation in the domain of military logistics tasks. However, because we are considering tasks that are performed in digital environments, the system may do more than just remind the user about the task context; it can actively save the context of interrupted tasks and restore it on demand. We expect that services of this type would significantly reduce cognitive load and minimize task resumption time.

Additionally, a complete rethinking of the metaphors used in interface design could reduce the problems related to context restore. In the current desktop interface, in order to complete a task – say, write a report – the user is forced to fragment the task into subtasks (such as using a word processor to write some text, collecting data from a spreadsheet, going back to the spreadsheet to insert the data). This fragmentation could be avoided by shifting from an *application-oriented* to a *task-oriented* approach to computer-based activities (Clauzel, Roda and Stojanov 2006; Gonzalez and Mark 2004; Kaptelinin and Czerwinski 2007; Roda, Stojanov and Clauzel 2006). The definition of real-task environments would also make it possible to evaluate their characteristics automatically and control for task interference, for example along the processing stages, perceptual modalities, visual channels and processing codes dimensions suggested by Wickens (2002) and discussed in section 2.2.2.1.

2.3.4 Disruption of primary task (distraction)

The type of breakdown situations discussed in the previous three sections all require directing the user's attention to a new primary task or to a related piece of information. This section analyses instead situations in which the objective of the system is not to help the user move from one primary task to another, but rather to help the user maintain awareness about secondary information whilst minimizing disruption of the primary task. As discussed in section 2.2.1, if task-irrelevant information is presented in a very conspicuous manner, bottom-up processes may cause enough activation to override the primary task, causing attention capture and thereby interrupting the primary task. A significant body of research reports on the negative effects of interruptions both on the effectiveness and on the agreeableness of task performance (Bailey, Konstan and Carlis 2001; Gillie and Broadbent 1989; Grundgeiger and Sanderson 2009; Zijlstra *et al.* 1999). These effects are modulated by several factors, including individual differences with respect to responses to task interruptions and restoration, the characteristics of the primary tasks, the characteristics of the interruption, and the context in which the primary task and interruption take place (Czerwinski, Horvitz and Wilhite 2004; Gievska, Lindeman and Sibert 2005; J. M. Hudson *et al.* 2002; McFarlane and Latorella 2002; Oulasvirta and Salovaara 2004; Speier, Vessey and Valacich 2003). The negative effects of interruption have also been reported to be more severe on mobile devices (Nagata 2003).

The problem of presenting information in a way that enables awareness but does not disrupt the primary task has been extensively studied within the *notification systems* literature. This literature covers several application domains, including messaging systems (Cutrell, Czerwinski and Horvitz 2001; Czerwinski, Cutrell and Horvitz 2000; Horvitz *et al.* 2003), alerting in military operations (Obermayer and Nugent 2000), shared document annotation (Brush *et al.* 2001), ambient displays (Altosaar *et al.* 2006), healthcare (Grundgeiger and Sanderson 2009), social awareness in collaborative activities (Carroll *et al.* 2003), end-user programming (Robertson *et al.* 2004), air traffic control (C.-Y. Ho *et al.* 2004; Ratwani *et al.* 2008) and many others. Frameworks for the evaluation of notification systems have been proposed in order to compare approaches and capitalize on design knowledge (Chewar, McCrickard and Sutcliffe 2004; McCrickard *et al.* 2003). McCrickard and his colleagues (McCrickard *et al.* 2003) propose to measure the effects of visual notification with respect to (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. Through a set of experiments evaluating notifications along the above parameters the authors were able to establish the fitness of specific notification mechanisms given the notification objective. For example, small-sized blast or fade-in-place animation were found to be best suited to goals of minimal attention reallocation (low interruption), immediate response (high reaction) and small knowledge gain (low comprehension) (see also McCrickard and Chewar 2003).

As this section discusses methodologies aimed at avoiding disruption of the primary task, we are particularly concerned with low interruption notifications. Ideally, in order to minimize disruption whilst ensuring that relevant content is appropriately attended to, notification systems should evaluate the relevance of the information to be delivered to the current user's context, and consequently select notification contents, timing and modality. One area of investigation that, to our knowledge, has not been explored is related to the evaluation of how different types of notification mechanisms may be affected by cross-talk effects, i.e., how the content of a notification message may interfere with the execution of the current task.

2.3.4.1 Interruption relevance Although experimental results report that notifications that are relevant to the user's task are less disruptive than irrelevant ones (Czerwinski, Cutrell and Horvitz 2000), the automatic evaluation of relevance is obviously not a trivial task because it requires semantic knowledge about user activity and about the interruption content. In order to address this problem, Arroyo Acosta (2007) has proposed to use semantic-knowledge-based systems which allow reasoning about concepts related to the user's goal and the interruption content. Another promising approach to relevance evaluation is metadata collection and analysis, as discussed by Schmitz *et al.* in chapter 8 of this volume.

Relevance evaluation remains, however, a very open field of research.

2.3.4.2 Notification contents An obvious design question in notification systems is how much should a notification say about the new task in order to minimize *interruption* whilst maximizing *reaction* and *comprehension*?

We distinguish three types of notification system. The first type, *pure notification*, is normally a fairly simple message providing a pointer to newly available information or tasks to be performed. Common examples include a 'jumping' icon pointing to newly available system updates, or a fading small window informing users of the arrival of email messages. The second type of notification system includes *awareness mechanisms*. Differently from pure notification, awareness mechanisms provide the information itself rather than a pointer to it. Examples include awareness displays in distributed collaborative systems and stock-monitoring systems. Finally, a notification can take the form of a *complete switch of context* (e.g., opening a new window with a new application). This last case can be considered as a notification with no task cueing or, to use McFarlane and Latorella's terminology, with no *annunciation signal* (McFarlane and Latorella 2002).

Appropriate content in pure notification mechanisms is necessary in order to supply users with enough information about the interrupting task so that they can make an informed decision on whether to redirect attention (this is one of the requirements of what Woods terms the pre-attentive reference (Woods 1995)). In other words, appropriate cueing enables *intentional dismissal* and *intentional integration* (McFarlane and Latorella 2002), which take place when the user is supplied with enough information to decide whether and when to interrupt the primary task (*intentional integration*) or to continue on it, disregarding the notification (*intentional dismissal*). A study in air traffic control environments (C.-Y. Ho *et al.* 2004) reports significant improvements in the management of interruptions with a notification system that provides users with information about the modality and timing of the prospective task. One obviously important element in the selection of notification content is ensuring that the message is informative enough whilst its comprehension can be accomplished with little disruption of the primary task. There is a continuous trade-off between providing enough information for intentional integration or dismissal and minimizing the chance of disrupting the primary task – see, for example, the discussion in Sarter 2005. This problem has been analysed, with respect to textual information, in the READY system. READY is a natural language interface that dynamically adapts to the user's time pressure and working-memory limitations by appropriately bundling messages to the user. For example, shorter messages are used if the user is under time pressure, longer ones if the user is more available (Bohnenberger *et al.* 2002; Jameson *et al.* 1999).

Awareness mechanisms are similar to pure notification signals in that they inform users about new information or pending tasks. However, awareness mechanisms may provide the user with large amounts of complex data because the core of the information is immediately made available. In order to support the extraction of relevant content from awareness displays, without disrupting performance of the primary task, Somervell and his colleagues (Somervell *et al.* 2002) propose to use peripheral visualization techniques. The authors argue that some visualization techniques, which have been shown to improve performance in situations of focused attention over large and/or complex amounts of data, could be used as peripheral mechanisms that, used under divided attention conditions, bring minimal disturbance to the primary task. In this manner, the benefits, in terms of high information comprehension, of visualization techniques can be integrated with the benefits, in terms of low interruption, of pure notification mechanisms. The use of visualization techniques to support users in situations of focused attention is briefly discussed in section 2.3.5.

2.3.4.3 Timing of interruption Notification timing impacts significantly on whether and how the interruption is perceived and on how much disruption it will bring to the current task. McFarlane and Latorella (2002: 5) propose four design solutions to schedule notifications:

> immediate, negotiated, mediated, and scheduled. Interruptions can be delivered at the soonest possible moment (immediate), or support can be given for the person to explicitly control when they will handle the interruption (negotiation). Another solution has an autonomous broker dynamically decide when best to interrupt the user (mediated), or to always hold all interruptions and deliver them at a prearranged time (scheduled)

and conclude that in most situations negotiation is the best choice: see Franke, Daniels and McFarlane 2002 and C.-Y. Ho *et al.* 2004 for examples of implementation of these strategies. The conclusions drawn by McFarlane and Latorella are in line with the prediction of the *goal-activation model* (see section 2.2.2.3) that the interval between an alert and the interruption proper is a critical period: the model in fact predicts that this time can be used to prepare to resume the interrupted task (Altmann and Trafton 2002).

More recent work has considered finer-grained analysis of interruption time, on the basis of either task knowledge or sensory input.

Task-knowledge-based timing relies on the analysis of the structure of the task being performed. Bailey and his colleagues (Bailey *et al.* 2006; Bailey and Konstan 2006; Iqbal *et al.* 2005) represent tasks as two-level hierarchies composed of coarse events which are 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. Czerwinski and her colleagues, for example, identify three task phases (planning, execution and evaluation) and analyse the different effects that interruptions have on these phases (Czerwinski, Cutrell and Horvitz 2000).

Sensory-input-based timing relies on sensors' input to detect user activity and the best times for interruption. On the basis of the observation that human beings can very efficiently, and in the presence of a very small number of cues, evaluate others' interruptibility, Fogarty and Hudson propose that interruptibility evaluation is attainable from simple sensors and that speech detectors are the most promising sensors (Fogarty *et al.* 2005; S. E. Hudson *et al.* 2003). Chen and Vertegaal (2004) 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, they derive the user's interruptibility.

Finally, our research group has successfully explored the integration of task knowledge and sensory input for the selection of the most appropriate interruption time by combining knowledge of a detailed task structure (Laukkanen, Roda and Molenaar 2007) with simple sensory input to evaluate the strength of breakpoints for possible interruptions (Molenaar and Roda 2008).

It should be noted that appropriate selection of interruption time is particularly critical in wireless devices because the user may be carrying/wearing such devices in a wide variety of situations. J. Ho and Intille (2005: 909) propose a context-aware mobile computing device that 'automatically detects postural and ambulatory activity transitions in real time using wireless accelerometers. This device was used to experimentally measure the receptivity to interruptions delivered at activity transitions, relative to those delivered at random times'. The authors conclude that messages are better received at times when the user is transitioning between different physical activities.

2.3.4.4 Interruption modality Interruption modality is an important factor determining how disruptive an interruption will be with respect to the primary task. In particular, multiple resource theories (see section 2.2.2) predict that cross-modal interruption presentation should generate lower disruption to the primary task. On the basis of this prediction, several cross-modality notification mechanisms have been designed and evaluated (Latorella 1998; Sarter 2006), reporting mixed results. The effects of modality, in fact, appear difficult to separate from other intervening effects.

It seems possible that the advantage obtained by certain notification mechanisms is due to automatic processing of the notification whilst, in other cases, it is due to a low interference between the primary task and the task of attending to the notification. In the latter case a task switch would intervene and therefore all costs involved in the task-switching process would contribute to the cost of processing the interruption (see section 2.2.2.3), including, in particular, the cost of returning to the primary task. As discussed in section 2.2.4.2, returning to the primary task is facilitated if an associative link is maintained with the primary task. In this situation, the advantages found in many experiments testing the effects of cross-modality notification could be ascribed to the maintenance of this link rather than to a multi-resource model (e.g., a spoken notification in a visual primary task environment would maintain the visibility of the primary task screen, thereby preserving the associative link). Support for this hypothesis has been found in several small-scale experiments (Field 1987; Ratwani *et al.* 2008).

2.3.5 Missing important events and information

Section 2.2.3.3 presents results showing that in certain situations an observer may miss very salient information. These events become increasingly common as our cognitive system is placed under greater strain due to larger or more complex perceptual input and/or increased task demands. Visualization techniques have been developed to address the needs of users working under these conditions. Results of the studies of pre-attentive and automatic processing (see section 2.2.3) supply a basis for providing information in a format that can be acquired by the user with minimal effort. Significant results have recently been achieved in addressing the question of how information can be presented so that the important elements stand out (Ware 2000), and the knowledge of specific psycho-physiological visual effects has significantly guided vision research. Healey and his colleagues (Healey, Booth and Enns 1996), for example, have shown that the use of pre-attentive features (hue and orientation) can support the discrimination of important information in the context of numerical estimation. Importantly, they highlight the fact that the use of features (especially for the representation of multi-dimensional data) should be guided by the rule of avoidance of feature interference. The authors have later extended their original work to analyse texture and colour. They show that colour discrimination is related to colour distance, linear separation and colour category, while texture is mainly discriminated by size and density (Healey and Enns 1999). Another feature used to support fast detection of visual stimuli is motion. Bartram and her colleagues (Bartram, Ware and Calvert 2003) show that icons with simple motions provide an effective notification mechanism which, especially in the periphery, is detected more easily than when the user is guided by colour and shape. Motion is reported as being effective in both the near and the far field of vision, further even small linear oscillation appears to be sufficient for discrimination and not to interfere with colour and shape coding (for a further discussion of peripheral displays supporting notification which also relates to the influence of time pressure, information density and secondary tasks, see Somervell *et al.* 2002). More recent research of perception in visualization has explored the possibilities opened up by the use of 3D displays. Kooi, in chapter 10 of this volume, demonstrates how, through the use of an extra real depth layer, it is possible to declutter the screen so that more information can be displayed whilst maintaining performance. A different approach is taken by Spence (2002), who proposes that scanning through a large amount of, possibly not organized, information can be made more efficient through Rapid Serial Visual Presentation which exploits the user's ability to recognize very quickly (in the range of 100 ms) whether some information is relevant.

Purely perceptual approaches to visualization, however, may not always be sufficient because, as discussed in section 2.2.3.3, what we see or do not see, as well as what is salient and to what degree, is dependent not only on properties of the external world, but also on our own internal state and objectives; and in particular on how we allocate attention. As a consequence, some authors move to what we call an *adaptive visualization approach* in which perceptual knowledge is integrated with knowledge related to the user's current interest, goals or focus of attention. Adaptive visualization is characterized by both the choice of perceptual techniques used for interaction and the method used to evaluate users' cognitive state. One method that has been extensively researched and applied is the use of users' gaze information as an indicator of users' attention allocation. This method is further discussed by Räihä, Hyrskykari and Majaranta in chapter 7 of this volume and a survey based on evaluation studies can be found in Toet 2006. A different method for the evaluation of the user's state is taken by Furnas (1999), who proposes to use a *degree of interest* function to evaluate which portion of the structure to be displayed is most likely to be of interest to the user. The degree of interest evaluation is applied to fisheye techniques for the display of a tree structure so that the degree of interest function aims at assigning an interest value to the tree nodes. Similar approaches include Card and Nation 2002 and Lamping and Rao 1994.

One could see search engines and ranking tools as an extreme example of adaptive visualization where the perceptual technique is extremely simple (just display relevant elements in order) and the degree of interest function is very complex and corresponds to the search algorithm. Search algorithms, however, do not take into account the past and current states of attention of the user. This could be achieved through the use of attention metadata as proposed by Schmitz and colleagues in chapter 8 of this volume.

In addition to the use of visualization techniques, the problem of users missing important events and information can be addressed due to the knowledge of cognitive effects such as change blindness. For example, particularly important changes that may be missed by the user can be made more obvious to the user through cues appearing in the focus area. To our knowledge, this technique has never been experimented with.

Pure visualization approaches are sometimes integrated with approaches relying on proactive system behaviour so that, rather than trying to attract the attention of a busy user to important events, these events are automatically treated by the system. Bosse and his colleagues (Bosse, van Maanen and Treur 2006), for example, propose a system capable of recognizing if the user is not attending to an important visual event (a track on screen representing a hostile aircraft), in which case the system takes over the task.

2.3.6 Habituation-related failures

There are two important aspects of automaticity as described in section 2.2.3 that should be considered in system design. First, automatic actions can be executed in the absence of attention and therefore without interfering with other tasks. Second, automatic actions are obligatory and therefore require an act of explicit control to be suppressed. These two observations reveal that whilst automaticity may enable the design of interfaces that, once learned, will place very little demand on cognitive processes, they will also engender situations in which actions are performed without awareness, and thus regardless of whether they are correct or not. Raskin (2000) notes, for example, that requiring a confirmation on file close is ineffective if the user always has to perform the same action to confirm. Assume that after a *file close* command a window appears in which one has to press the return key to confirm; because, in the great majority of cases, one would press the return key to confirm, the sequence *file close followed by press return* will become a habitual action and it will be executed also in the rare cases when the confirmation shouldn't have been given. The consequence is that 'any confirmation step that elicits any fixed response soon becomes useless' (Raskin 2000: 22). On the other hand, if the interface is appropriately designed its use may become automatic, with all the advantages of automatic performance discussed earlier in this chapter.

In section 2.2.3.2 we saw that learned automaticity does not transfer well; this implies that habituation is better achieved in interfaces that, given a sequence of commands, always produce the same behaviour. This constraint may interfere with designs that strive dynamically to adapt the interface to the user's need, or simply to the context of the operation. To address the latter case, Raskin proposes that operation through *modes* should be replaced by operation through *quasi-modes*. Modes are contexts in which a user action may result in different behaviours of the system. A simple example is the different behaviour the system has when typing, depending on whether the Caps Lock key is engaged. The Caps Lock key defines two modes of operation of the system. The problem of modes is that users may perform actions without being aware of the mode they are in and therefore trigger unexpected system behaviours (unawareness of the Caps Lock mode, for example, may hinder the work of even the most experienced typist). Raskin argues that quasi-modes, such as the holding of the Shift key while typing, integrate the operational context in the gesture required for the user to perform the action. In this manner the interface can operate in a context-dependent manner whilst the user can acquire action automaticity: see Raskin (2000: chapter 3) for a thorough discussion.

2.4 Other relevant research areas

Although this chapter covers a wide variety of aspects of attention, some other aspects are, at least potentially, relevant to the design of attention-aware systems. The first of these is the relation between attention and emotion. Studies in this area analyse the effects of emotional stimuli depending on their valence, including reactions to happy/sad, fearful/safe and generally positive/negative stimuli. Emotional valence appears to be evaluated both at perceptual and at semantic level; it therefore influences reaction to stimuli at all levels of processing (see Compton 2003 for a review, and Lavie 2005, Lim, Padmala and Pessoa 2008 and Pêcher, Lemercier and Cellier 2009 for examples of relevant more recent work). Interfaces that strive to manage emotional aspects include those based on animated agents. Chapter 6 of this book discusses the impact of one such interface on the support of attention allocation.

Second, findings in the area of social attention could improve our understanding of the constraints and motivations guiding attention allocation. Two aspects are particularly relevant. On the one hand there is the research covered in the field of *joint attention* (Eilan *et al.* 2005; Frischen, Bayliss and Tipper 2007; Moore and Dunham 1995), which is at the heart of communication and collaboration. On the other hand, there is the research covered in the field of *collective attention*, which addresses questions such as: How do communities of people allocate attention to a given item? What processes underlie collective attention? How can these processes be supported? Some of these aspects are treated in Huberman 2008 and in chapters 8 and 12 of this book.

Finally, several studies have addressed the question of how attention allocation may vary amongst individuals; this research goes under the name of *individual differences*. In particular, age and gender appear to be significant factors (Engle, Kane and Tuholski 1999; Frischen, Bayliss and Tipper 2007; Lavie 2005). It is conceivable that interfaces capable of adapting to these individual differences may provide for a better attention support.

2.5 Conclusions

Many years of studies in cognitive psychology and, more recently, in neuroscience have built a body of knowledge about the central role of attention in human physical and mental activity. Although theories concerning various aspects of attention are sometimes scattered and controversial, many results are very relevant to the design of attention-aware systems. In this chapter we have discussed how such results help us design systems that address the problem of *attentional breakdowns*. For example, research in the area of divided attention indicates that the time span between a notification and the actual task switch, together with the associative cues provided by the system in this period, have a significant impact on task resumption; research in perception and automatic processes has provided guidelines for the design of efficient visualization mechanisms; multiple resources theories provide a path to minimize task interference in multitasking environments; and there are many other such results.

Whilst some of these results are directly applicable to system design, many others require further research in order to evaluate their effects in real-world environments where the number of perceptual stimuli and the cognitive state of the user cannot be controlled as they are in laboratory settings.

2.6 Acknowledgements

This research was partially supported by a grant from the European Commission under the FP6 Framework project AtGentive IST 4–027529-STP.

2.7 References

Allport, A. 1993. Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years, in D. E. Meyer and S. Kornblum (eds.), *Attention and Performance*. London: MIT Press: vol. XIV: 183–218.

Altmann, E. M., and Gray, W. D. 2000. Managing attention by preparing to forget. Paper presented at the IEA 2000/HFES 2000 Congress. San Diego: 152–5

Altmann, E. M., and Trafton, J. G. 2002. Memory for goals: An activation-based model, *Cognitive Science* 26(1): 39–83

Altosaar, M., Vertegaal, R., Sohn, C., and Cheng, D. 2006. AuraOrb: Social notification appliance. Paper presented at the CHI '06, *Extended Abstracts on Human Factors in Computing Systems*. New York: 381–6

Anderson, J. R. 1983. A spreading activation theory of memory, *Journal of Verbal Learning and Verbal Behavior* 22(3): 261–95

1996. ACT: A simple theory of complex cognition, *American Psychologist* 51(4): 355–65

Arroyo Acosta, E. 2007. Mediating disruption in human–computer interaction from implicit metrics of attention. Doctoral thesis, Massachusetts Institute of Technology, Cambridge, MA

Awh, E., Vogel, E. K., and Oh, S. H. 2006. Interactions between attention and working memory, *Neuroscience* 139(1): 201–8

Baddeley, A. 1986. *Working Memory*. Oxford: Oxford University Press

2003. Working memory: Looking back and looking forward, *Nature Reviews Neuroscience* 4: 29–39

Baddeley, A., and Hitch, G. J. 1974. Working memory, in G. A. Bower (ed.), *Recent Advances in Learning and Motivation*. New York: Academic Press: vol. VIII: 47–89

Baddeley, A., and Weiskrantz, L. (eds.) 1993. *Attention: Selection, awareness and control. A tribute to Donald Broadbent*. Oxford: Clarendon Press

Bailey, B. P., Adamczyk, P. D., Chang, T. Y., and Chilson, N. A. 2006. A framework for specifying and monitoring user tasks, *Computers in Human Behavior* 22(4): 709–32

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

Bailey, B. P., Konstan, J. A., and Carlis, J. V. 2001. The effects of interruptions on task performance, annoyance, and anxiety in the user interface. Paper presented at the INTERACT '01: 593–601

Bargh, J. A., and Chartrand, T. L. 1999. The unbearable automaticity of being, *American Psychologist* 54(7): 462–79

2000. The mind in the middle: A practical guide to priming and automaticity research, in H. T. Reis and C. M. Judd (eds.), *Handbook of Research Methods in Social and Personality Psychology*. New York: Cambridge University Press: vol. XII: 253–85

Bartram, L., Ware, C., and Calvert, T. 2003. Moticons: Detection, distraction and task, *International Journal of Human–Computer Studies* 58(5): 515–45

Bichot, N. P., Cave, K. R., and Pashler, H. 1999. Visual selection mediated by location: Feature-based selection of noncontiguous locations, *Perception and Psychophysics* 61(3): 403–23

Bohnenberger, T., Brandherm, B., Grossmann-Hutter, B., Heckmann, D., and Wittig, F. 2002. Empirically grounded decision-theoretic adaptation to situation-dependent resource limitations, *Künstliche Intelligenz* 3: 10–16

Bosse, T., van Maanen, P.-P., and Treur, J. 2006. A cognitive model for visual attention and its application. Paper presented at the International Conference on Intelligent Agent Technology. IAT '06. IEEE/WIC/ACM 255–62

Boudreau, C. E., Williford, T. H., and Maunsell, J. H. R. 2006. Effects of task difficulty and target likelihood in area V4 of macaque monkeys, *Journal of Neurophysiology* 96: 2377–87

Broadbent, D. E. 1958. *Perception and Communication*. London: Pergamon Press

Brush, B., Bargeron, D., Gupta, A., and Grudin, J. 2001. *Notification for Shared Annotation of Digital Documents* (Technical Report No. MSR-TR-2001–87): Microsoft Research

Buehner, M., Krumm, S., Ziegler, M., and Pluecken, T. 2006. Cognitive abilities and their interplay: Reasoning, crystallized intelligence, working memory components, and sustained attention, *Journal of Individual Differences* 27(2): 57–72

Burgess, P. W., and Shallice, T. 1997. The relationship between prospective and retrospective memory: Neuropsychological evidence, in M. A. Conway (ed.), *Cognitive Models of Memory*. Cambridge, MA: MIT Press: 247–72

Card, S. K., and Nation, D. 2002. Degree-of-interest trees: A component of an attention–reactive user interface. Paper presented at the Advanced Visual Interfaces Conference: 231–45

Carroll, J. M., Neale, D. C., Isenhour, P. L., Rosson, M. B., and McCrickard, D. S. 2003. Notification and awareness: Synchronizing task-oriented collaborative activity, *International Journal of Human–Computer Studies* 58(5): 605–32

Cave, K. R., and Bichot, N. P. 1999. Visuospatial attention: Beyond a spotlight model, *Psychonomic Bulletin and Review* 6(2): 204–23

Chen, D., and 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*. Vienna, Austria: ACM Press: 1513–16

Chewar, C. M., McCrickard, D. S., and Sutcliffe, A. G. 2004. Unpacking critical parameters for interface design: Evaluating notification systems with the IRC framework. Paper presented at the Conference on Designing Interactive Systems: Processes, Practices, Methods and Techniques (DIS 2004), Slovakia: 279–88

Chun, M. M., and Wolfe, J. 2001. Visual attention, in E. B. Goldstein (ed.), *Blackwell's Handbook of Perception*. Oxford: Blackwell: 272–310

Clauzel, D., Roda, C., and Stojanov, G. 2006. Tracking task context to support resumption. Paper presented at the HCI 2006 workshop on computer assisted recording, pre-processing, and analysis of user interaction data, London: 43–54

Colom, R., Rebollo, I., Abad, F. J., and Shih, P. C. 2006. Complex span tasks, simple span tasks, and cognitive abilities: A reanalysis of key studies, *Memory and Cognition* 34(1): 158–71

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): 115–29

Conway, A. R. A., Cowan, N., Bunting, M. F., Therriault, D. J., and Minkoff, S. R. B. 2002. A latent variable analysis of working memory capacity, short-term memory capacity, processing speed, and general fluid intelligence, *Intelligence* 30: 163–84

Conway, A. R. A., Moore, A. B., and Kane, M. J. 2009. Recent trends in the cognitive neuroscience of working memory, *Cortex* 45(2): 262–8

Cowan, N. 1988. Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system, *Psychological Bulletin* 104(2): 163–91

Cowan, N., Morey, C. C., Chen, Z., Gilchrist, A. L., and Saults, J. S. 2008. Theory and measurement of working memory capacity limits, *The Psychology of Learning and Motivation* 49: 49–104

Cutrell, E., Czerwinski, M., and Horvitz, E. 2001. Notification, disruption, and memory: Effects of messaging interruptions on memory and performance. Paper presented at the Interact 2001 IFIP Conference on Human–Computer Interaction, Tokyo: 263–9

Czerwinski, M., Cutrell, E., and Horvitz, E. 2000. Instant messaging: Effects of relevance and time. Paper presented at the HCI 2000 – 14th British HCI Group Annual Conference: 71–6

Czerwinski, M., and Horvitz, E. 2002. An investigation of memory for daily computing events. Paper presented at the HCI 2002: 230–45

Czerwinski, M., Horvitz, E., and Wilhite, S. 2004. A diary study of task switching and interruptions. Paper presented at the SIGCHI Conference on Human Factors in Computing Systems. Vienna: 175–82

Desimone, R., and Duncan, J. 1995. Neural mechanisms of selective visual attention, *Annual Review of Neuroscience* 18: 193–222

Diamond, A. 2005. Attention-deficit disorder (attention-deficit/hyperactivity disorder without hyperactivity), A neurobiologically and behaviorally distinct disorder from attention-deficit/hyperactivity disorder (with hyperactivity), *Development and Psychopathology* 17: 807–25

Dijksterhuis, A., Bos, M. W., Nordgren, L. F., and van Baaren, R. B. 2006. On making the right choice: The deliberation-without-attention effect, *Science* 311: 1005–7

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

Driver, J., Davis, G., Russell, C., Turatto, M., and Freeman, E. 2001. Segmentation, attention and phenomenal visual objects, *Cognition* 80(1–2): 61–95

Duncan, J., and Humphreys, G. W. 1989. Visual search and stimulus similarity, *Psychological Review* 96(3): 433–58

Duncan, J., Martens, S., and Ward, R. 1997. Restricted attentional capacity within but not between sensory modalities, *Nature* 387(6635): 808–10

Eichenbaum, H. 1997. Declarative memory: Insights from cognitive neurobiology, *Annual Review of Psychology* 48(1): 547–72

Eilan, N., Hoerl, C., McCormack, T., and Roessler, J. (eds.) 2005. *Joint Attention: Communication and Other Minds.* Oxford: Oxford University Press

Einstein, G. O., and McDaniel, M. A. 1996. Retrieval processes in prospective memory: Theoretical approaches and some new empirical findings, in M. Brandimonte, G. O. Einstein and M. A. McDaniel (eds.), *Prospective Memory: Theory and Applications*. Mahwah, NJ: Lawrence Erlbaum Associates: 115–42

Eisenhauer, M., Lorenz, A., Zimmermann, A., Duong, T. V., and James, F. 2005. Interaction by movement – One giant leap for natural interaction in mobile guides. Paper presented at the mobileHCI 05, Salzburg: www.comp.lancs.ac.uk/~kc/mguides05/pdfs/Eisenhauer_InteractionMovement.pdf (retrieved July 2010)

Eldridge, M., Sellen, A., and Bekerian, D. 1992. *Memory Problems at Work: Their Range, Frequency, and Severity* (Technical Report EPC-1992-129) Rank Xerox Research Centre

Engle, R. W. 2002. Working memory capacity as executive attention, *Current Directions in Psychological Science* 11(1): 19–23

Engle, R. W., Kane, M. J., and Tuholski, S. W. 1999. Individual differences in working memory capacity and what they tell us about controlled attention, general fluid intelligence, and functions of the prefrontal cortex, in A. Miyake and P. Shah (eds.), *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control.* New York: Cambridge University Press: 102–34

Engle, R. W., Tuholski, S. W., Laughlin, J. E., and Conway, A. R. A. 1999. Working memory, short-term memory, and general fluid intelligence: A latent-variable approach, *Journal of Experimental Psychology: General* 128(3): 309–31

Fahy, F. L., Riches, I. P., and Brown, M. W. 1993. Neuronal activity related to visual recognition memory: Long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex, *Experimental Brain Research* 96(3): 457–72

Fernandez-Duque, D., and Johnson, M. 2002. Cause and effect theories of attention: The role of conceptual metaphors, *Review of General Psychology* 6(2): 153–65

Field, G. E. 1987. Experimentus interruptus, *SIGCHI Bulletin* 19(2): 42–6

Fogarty, J., Hudson, S. E., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., *et al*. 2005. Predicting human interruptibility with sensors, *ACM Transactions on Computer–Human Interaction* 12(1): 119–46

Franconeri, S. L., and Simons, D. J. 2003. Moving and looming stimuli capture attention, *Perception and Psychophysics* 65(7): 999–1010

Franke, J. L., Daniels, J. J., and McFarlane, D. C. 2002. Recovering context after interruption. Paper presented at the 24th Annual Meeting of the Cognitive Science Society (CogSci 2002): 310–15

Freeman, E., and Gelernter, D. 2007. Beyond lifestreams: The inevitable demise of the desktop metaphor, in Kaptelinin and Czerwinski: 19–48

Frischen, A., Bayliss, A. P., and Tipper, S. P. 2007. Gaze cueing of attention: Visual attention, social cognition, and individual differences, *Psychological Bulletin* 133(4): 694–724

Furnas, G. W. 1999. The FISHEYE view: A new look at structured files, in S. K. Card, J. Mackinlay and B. Shneiderman (eds.), *Readings in Information Visualization: Using Vision to Think*. San Francisco: Morgan Kaufmann: 312–30

Gibson, B. S., Folk, C., Theeuwes, J., and Kingstone, A. (eds.) 2008. Attentional capture, a special issue of *Visual Cognition* 16(2/3): Psychology Press

Gibson, B. S., and Peterson, M. A. 2001. Inattentional blindness and attentional capture: Evidence for attention-based theories of visual salience, in C. L. Folk and B. S. Gibson (eds.), *Attraction, Distraction and Action: Multiple Perspectives on Attentional Capture*. New York: Elsevier Science: 51–76

Gievska, S., Lindeman, R., and Sibert, J. 2005. Examining the qualitative gains of mediating human interruptions during HCI. Paper presented at the 11th International Conference on Human–Computer Interaction (HCI International 2005). Mahwah, NJ: 605–14

Gillie, T., and Broadbent, D. 1989. What makes interruptions disruptive? A study of length, similarity, and complexity, *Psychological Research* 50(4): 243–50

Goldman-Rakic, P. S. 1987. Circuitry of the prefrontal cortex and the regulation of behavior by representational memory, in V. B. Mountcastle, F. Plum and S. R. Geiger (eds.), *Handbook of Neurobiology*. Bethesda, MD: American Physiological Society: 373–417

Gonzalez, V. M., and Mark, G. 2004. 'Constant, constant, multi-tasking craziness': Managing multiple working spheres. Paper presented at the SIGCHI Conference on Human Factors in Computing Systems. Vienna: 113–20

Grundgeiger, T., and Sanderson, P. 2009. Interruptions in healthcare: Theoretical views, *International Journal of Medical Informatics* 78(5): 293–307

Healey, C. G., Booth, K. S., and Enns, J. T. 1996. High-speed visual estimation using preattentive processing, *ACM Transactions on Computer–Human Interaction* 3(2): 107–35

Healey, C. G., and Enns, J. T. 1999. Large datasets at a glance: Combining textures and colors in scientific visualization, *IEEE Transactions on Visualization and Computer Graphics* 5(2): 145–67

Hillyard, S. A., Vogel, E. K., and Luck, S. J. 1998. Sensory gain control (amplification) as a mechanism of selective attention: Electro-physiological and neuroimaging evidence, *Philosophical Transactions: Biological Sciences* 353(1373): 1257–70

Ho, C., Tan, H. Z., and Spence, C. 2005. Using spatial vibrotactile cues to direct visual attention in driving scenes, *Transportation Research Part F: Traffic Psychology and Behaviour* 8(6): 397–412

Ho, C.-Y., Nikolic, M. I., Waters, M. J., and Sarter, N. B. 2004. Not now! Supporting interruption management by indicating the modality and urgency of pending tasks, *Human Factors: The Journal of the Human Factors and Ergonomics Society* 46: 399–409

Ho, J., and Intille, S. S. 2005. Using context-aware computing to reduce the perceived burden of interruptions from mobile devices. Paper presented at the CHI 2005. Portland, OR: 909–18

Hommel, B. 2010. Grounding attention in action control: The intentional control of selection, in B. J. Bruya (ed.), *Effortless Attention: A New Perspective in the Cognitive Science of Attention and Action*. Cambridge, MA: MIT Press: 121–40

Hommel, B., Ridderinkhof, K. R., and Theeuwes, J. 2002. Cognitive control of attention and action: Issues and trends, *Psychological Research* 66(4): 215–19

Horvitz, E., Kadie, C., Paek, T., and Hovel, D. 2003. Models of attention in computing and communication: From principles to applications, *Communications of the ACM* 46(3): 52–9

Huberman, B. A. 2008. Crowdsourcing and attention, *Computer* 41(11): 103–5

Hudson, J. M., Christensen, J., Kellogg, W. A., and Erickson, T. 2002. 'I'd be overwhelmed, but it's just one more thing to do': Availability and interruption in research management. Paper presented at the SIGCHI Conference on Human Factors in Computing Systems. Minneapolis: 97–104

Hudson, S. E., Fogarty, J., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., *et al.* 2003. Predicting human interruptibility with sensors: A Wizard of Oz feasibility study. Paper presented at the CHI 2003. Fort Lauderdale, FL: 257–64

Hussain, F., and Wood, S. 2009. Modelling the efficiencies and interactions of attentional networks, in L. Paletta and J. K. Tsotsos (eds.), *Attention in Cognitive Systems*. Berlin: Springer-Verlag: 139–52

Iqbal, S. T., Adamczyk, P. D., Zheng, X. S., and Bailey, B. P. 2005. Towards an index of opportunity: Understanding changes in mental workload during task execution. Paper presented at the SIGCHI Conference on Human Factors in Computing Systems. Portland, OR: 311–20

Itti, L. 2005. Models of bottom-up attention and saliency, in L. Itti, G. Rees and J. K. Tsotsos (eds.), *Neurobiology of Attention*. San Diego: Elsevier: 576–82

Itti, L., Koch, C., and Niebur, E. 1998. A model of saliency-based visual attention for rapid scene analysis, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29: 1254–9

Jameson, A., Schafer, R., Weis, T., Berthold, A., and Weyrath, T. 1999. Making systems sensitive to the user's time and working memory constraints, in *Proceedings of the 4th International Conference on Intelligent User Interfaces*. Los Angeles: ACM Press: 79–86

Kahneman, D. 1973. *Attention and Effort*. Englewood Cliffs, NJ: Prentice Hall

Kaptelinin, V., and Czerwinski, M. (eds.) 2007. *Beyond the Desktop Metaphor: Designing Integrated Digital Work Environments*. Cambridge, MA: MIT Press

Kardiasmenos, K. S., Clawson, D. M., Wilken, J. A., and Wallin, M. T. 2008. Prospective memory and the efficacy of a memory strategy in multiple sclerosis, *Neuropsychology* 22(6): 746–54

Karger, D. R. 2007. Haystack: Per-user information environment based on semistructured data, in Kaptelinin and Czerwinski: 49–100

Kliegel, M., Mackinlay, R., and Jäger, T. 2008. Complex prospective memory: Development across the lifespan and the role of task interruption, *Developmental Psychology* 44(2): 612–17

Kliegel, M., and Martin, M. 2003. Prospective memory research: Why is it relevant? *International Journal of Psychology* 38(4): 193–4

Koch, C., and Tsuchiya, N. 2007. Attention and consciousness: Two distinct brain processes, *Trends in Cognitive Sciences* 11(1): 16–22

Kvavilashvili, L., Messer, D. J., and Ebdon, P. 2001. Prospective memory in children: The effects of age and task interruption, *Developmental Psychology* 37(3): 418–30

LaBerge, D. 2002. Attentional control: Brief and prolonged, *Psychological Research* 66(4): 220–33

Lamming, M., Brown, P., Carter, K., Eldridge, M., Flynn, M., Louie, G., *et al.* 1994. The design of a human memory prosthesis, *Computer Journal* 37(3): 153–63

Lamming, M., and Flynn, M. 1994. 'Forget-me-not': Intimate computing in support of human memory. Paper presented at the FRIEND21 Symposium on Next Generation Human Interfaces, Tokyo. Also available as RXRC TR 94–103, 61 Regent St, Cambridge, UK

Lamping, J., and Rao, R. 1994. Laying out and visualizing large trees using a hyperbolic space. Paper presented at the 7th Annual ACM Symposium on User Interface Software and Technology. Marina del Rey, CA

Latorella, K. A. 1998. Effects of modality on interrupted flight deck performance: Implications for data link. Paper presented at the Human Factors and Ergonomics Society 42nd Annual Meeting, Santa Monica, CA: 87–91

Laukkanen, J., Roda, C., and Molenaar, I. 2007. Modelling tasks: A requirements analysis based on attention support services. Paper presented at the workshop on contextualized attention metadata: personalized access to digital resources, CAMA 2007 at the ACM IEEE Joint Conference on Digital Libraries, 17–23 June 2007, Vancouver: 1–7

Lavie, N. 2005. Distracted and confused? Selective attention under load, *Trends in Cognitive Sciences* 9: 75–82

Lavie, N., Hirst, A., de Fockert, J. W., and Viding, E. 2004. Load theory of selective attention and cognitive control, *Journal of Experimental Psychology: General* 133(3): 339–54

Lavie, N., and Tsal, Y. 1994. Perceptual load as a major determinant of the locus of selection in visual attention, *Perception and Psychophysics* 56(2): 183–97

Lim, S.-L., Padmala, S., and Pessoa, L. 2008. Affective learning modulates spatial competition during low-load conditions, *Neuropsychologia* 46(5): 1267–78

Logan, G. D. 1988. Toward an instance theory of automatization, *Psychological Review* 95(4): 492–527

1992. Attention and preattention in theories of automaticity, *American Journal of Psychology* 105(2): 317–39

Mack, A. 2003. Inattentional blindness: Looking without seeing, *Current Directions in Psychological Science* 12(5): 180–4

Mack, A., and Rock, I. 1998. *Inattentional Blindness*. Cambridge, MA: MIT Press

Mark, G., Gonzalez, V. M., and Harris, J. 2005. No task left behind? Examining the nature of fragmented work. Paper presented at the SIGCHI conference on Human Factors in Computing Systems. Portland, OR: 321–30

Marsh, R. L., Hicks, J. L., and Bryan, E. S. 1999. The activation of un-related and canceled intentions, *Memory and Cognition* 27(2): 320–7

Mase, K., Sumi, Y., and Fels, S. 2007. Welcome to the special issue on memory and sharing of experience for the *Journal of Personal and Ubiquitous Computing*, *Journal of Personal and Ubiquitous Computing* 11(4): 213–14

McCrickard, D. S., Catrambone, R., Chewar, C. M., and 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): 547–82

McCrickard, D. S., and Chewar, C. M. 2003. Attuning notification design to user goals and attention costs, *Communications of the ACM* 46(3): 67–72

McDaniel, M. A., Guynn, M. J., Einstein, G. O., and Breneiser, J. 2004. Cue-focused and reflexive-associative processes in prospective memory retrieval, *Journal of Experimental Psychology: Learning, Memory, and Cognition* 30(3): 605–14

McFarlane, D. C., and Latorella, K. A. 2002. The scope and importance of human interruption in human–computer interaction design, *Human–Computer Interaction* 17(1): 1–62

McMains, S. A., and Somers, D. C. 2004. Multiple spotlights of attentional selection in human visual cortex, *Neuron* 42: 677–86

McNab, F., and Klingberg, T. 2008. Prefrontal cortex and basal ganglia control access to working memory, *Nature Neuroscience* 11(1): 103–7

Meiran, N., Chorev, Z., and Sapir, A. 2000. Component processes in task switching, *Cognitive Psychology* 41(3): 211–53

Miller, G. A. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information, *Psychological Review* 63: 81–97

Molenaar, I., and Roda, C. 2008. Attention management for dynamic and adaptive scaffolding, *Pragmatics and Cognition* 16(2): 225–71

Monsell, S. 2003. Task switching, *Trends in Cognitive Sciences* 7(3): 134–40

Moore, C., and Dunham, P. J. (eds.) 1995. *Joint Attention: Its Origins and Role in Development*. Hillsdale, NJ: Lawrence Erlbaum Associates

Nagata, S. F. 2003. Multitasking and interruptions during mobile web tasks. Paper presented at the 47th Annual Meeting of the Human Factors and Ergonomics Society: 1341–5

Navalpakkam, V., and Itti, L. 2006. An integrated model of top-down and bottom-up attention for optimal object detection. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR): 2049–56

Navon, D., and Miller, J. 1987. Role of outcome conflict in dual-task interference, *Journal of Experimental Psychology: Human Perception and Performance* 13: 435–48

Nisbett, R. E., and Masuda, T. 2003. Culture and point of view, *Proceedings of the National Academy of Sciences of the United States of America* 100(19): 11163–70

Norman, D., and Shallice, T. 1986. Attention to action: Willed and automatic control of behavior, in R. J. Davidson, G. E. Schwartz and D. Shapiro (eds.), *Consciousness and Self-Regulation*. New York: Plenum: vol. IV: 1–18

Oberauer, K. 2002. Access to information in working memory: Exploring the focus of attention, *Journal of Experimental Psychology: Learning, Memory, and Cognition* 28(3): 411–21

Obermayer, R. W., and Nugent, W. A. 2000. Human–computer interaction for alert warning and attention allocation systems of the multi-modal watchstation. Paper presented at the SPIE 2000 – The International Society for Optical Engineering: 14–22

O'Conaill, B., and Frohlich, D. 1995. Timespace in the workplace: Dealing with interruptions. Paper presented at the CHI '95 Conference, Denver, CO: 262–3

Olivers, C. N. L., and Nieuwenhuis, S. 2005. The beneficial effect of concurrent task-irrelevant mental activity on temporal attention, *Psychological Science* 16: 265–9

O'Regan, J. K. 1992. Solving the 'real' mysteries of visual perception: The world as an outside memory, *Canadian Journal of Psychology* 46(3): 461–88

O'Regan, J. K., and Noë, A. 2001. A sensorimotor account of vision and visual consciousness, *Behavioral and Brain Sciences* 24(5): 939–1031

Oulasvirta, A. 2004. Task demands and memory in web interaction: A levels-of-processing approach, *Interacting with Computers* 16(2): 217–41

Oulasvirta, A., Kärkkäinen, L., and Laarni, J. 2005. Expectations and memory in link search, *Computers in Human Behavior* 21(5): 773–89

Oulasvirta, A., and Salovaara, A. 2004. A cognitive meta-analysis of design approaches to interruptions in intelligent environments, in *Extended Abstracts of the 2004 Conference on Human Factors and Computing Systems*. Vienna: ACM Press: 1155–8

Parasuraman, R., and Davis, D. R. (eds.) 1984. *Varieties of Attention*. San Diego: Academic Press

Pashler, H. 1994. Dual-task interference in simple tasks: Data and theory, *Psychological Bulletin* 116(2): 220–44

2000. Task switching and multitask performance, in S. Monsell and J. Driver (eds.), *Attention and Performance XVIII: Control of Mental Processes*. Cambridge, MA: MIT Press: 277–308

Pashler, H., and Johnston, J. C. 1998. Attentional limitations in dual-task performance, in H. Pashler (ed.), *Attention*. Hove: Psychology Press: 155–189

Pêcher, C., Lemercier, C., and Cellier, J.-M. 2009. Emotions drive attention: Effects on driver's behaviour. *Safety Science* 47(9): 1254–9

Polster, M. R., Nadel, L., and Schacter, D. L. 2007. Cognitive neuroscience analyses of memory: A historical perspective, *Journal of Cognitive Neuroscience* 3(2): 95–116

Posner, M. I. 1982. Cumulative development of attention theory, *American Psychologist* 37: 168–79

1994. Attention: The mechanisms of consciousness, *Proceedings of the National Academy of Sciences of the United States of America* 91(16): 7398–403

Posner, M. I., and Fan, J. 2007. Attention as an organ system, in J. Pomerantz (ed.), *Neurobiology of Perception and Communication: From Synapse to Society. The IVth De Lange Conference*. Cambridge, UK: Cambridge University Press

Posner, M. I., and Rothbart, M. K. 2007. Research on attention networks as a model for the integration of psychological science, *Annual Review of Psychology* 58(1): 1–23

Postle, B. R. 2006. Working memory as an emergent property of the mind and brain, *Neuroscience* 139(1): 23–38

Raskin, J. 2000. *The Humane Interface: New Directions for Designing Interactive Systems*. New York: Addison-Wesley

Ratwani, R. M., Andrews, A. E., Sousk, J. D., and Trafton, J. G. 2008. The effect of interruption modality on primary task resumption, *Human Factors and Ergonomics Society Annual Meeting Proceedings* 52: 393–7

Rensink, R. A. 2007. The modeling and control of visual perception, in W. Gray (ed.), *Integrated Models of Cognitive Systems*. New York: Oxford University Press: 132–48

Robertson, T. J., Prabhakararao, S., Burnett, M., Cook, C., Ruthruff, J. R., Beckwith, L., *et al.* 2004. Impact of interruption style on end-user debugging. Paper presented at the SIGCHI Conference on Human Factors in Computing Systems. Vienna: 287–94

Roda, C., and Nabeth, T. 2007. Supporting attention in learning environments: Attention support services and information management. Paper presented at

Creating New Experiences on a Global Scale. Second European Conference on Technology Enhanced Learning, EC-TEL 2007, Crete: 277–91

Roda, C., Stojanov, G., and Clauzel, D. 2006. Mind-prosthesis metaphor for design of human–computer interfaces that support better attention management. Paper presented at the AAAI 2006 Fall Symposium on 'Interaction and Emergent Phenomena in Societies of Agents', Arlington, VA: 52–9

Roda, C., and Thomas, J. 2006. Attention aware systems: Theories, applications, and research agenda, *Computers in Human Behavior* 22(4): 557–87

Rogers, R., and Monsell, S. 1995. Costs of a predictable switch between simple cognitive tasks, *Journal of Experimental Psychology: General* 124(2): 207–31

Sarter, N. 2005. Graded and multimodal interruption cueing in support of preattentive reference and attention management, *Human Factors and Ergonomics Society Annual Meeting Proceedings* 49: 478–81

2006. Multiple-resource theory as a basis for multimodal interface design: Success stories, qualifications, and research needs, in A. F. Kramer, D. A. Wiegmann and A. Kirlik (eds.), *Attention: From Theory to Practice*. Oxford: Oxford University Press: 187–93

Schacter, D. L. 1987. Implicit memory: History and current status, *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13(3): 501–18

Schill, K., Zetzsche, C., and Hois, J. 2009. A belief-based architecture for scene analysis: From sensorimotor features to knowledge and ontology, *Fuzzy Sets and Systems* 160(10): 1507–16

Scholl, B. J. 2001. Objects and attention: The state of the art, *Cognition* 80(1–2): 1–46

Sellen, A. J., Louie, G., Harris, J. E., and Wilkins, A. J. 1997. What brings intentions to mind? An in situ study of prospective memory, *Memory* 5(4): 483–507

Simons, D. J. Visual cognition lab – demos and stimuli. Retrieved June 2009 from: http://viscog.beckman.illinois.edu/djs_lab/demos.html.

Simons, D. J., and Chabris, C. F. 1999. Gorillas in our midst: Sustained inattentional blindness for dynamic events, *Perception* 28(9): 1059–74

Simons, D. J., and Rensink, R. A. 2005. Change blindness: past, present, and future, *Trends in Cognitive Sciences* 9(1): 16–20

Somervell, J., McCrickard, D. S., North, C., and Shukla, M. 2002. An evaluation of information visualization in attention-limited environments. Paper presented at the Symposium on Data Visualisation 2002, Barcelona

Speier, C., Vessey, I., and Valacich, J. S. 2003. The effects of interruptions, task complexity, and information presentation on computer-supported decision-making performance, *Decision Sciences* 34(4): 771–97

Spence, R. 2002. Rapid, serial and visual: A presentation technique with potential, *Information Visualization* 1(1): 13–19

Stinson, C. 2009. Searching for the source of executive attention, *Psyche* 15(1): 137–54

Takegata, R., Brattico, E., Tervaniemi, M., Varyagina, O., Näätänen, R., and Winkler, I. N. 2005. Preattentive representation of feature conjunctions for concurrent spatially distributed auditory objects, *Cognitive Brain Research* 25(1): 169–79

Thorpe, S., Fize, D., and Marlot, C. 1996. Speed of processing in the human visual system, *Nature* (381): 520–2

Toet, A. 2006. Gaze directed displays as an enabling technology for attention aware systems, *Computers in Human Behavior* 22(4): 615–47

Treisman, A. 1960. Contextual cues in selective listening, *Quarterly Journal of Experimental Psychology* 12: 242–8

1985. Preattentive processing in vision, *Computer Vision, Graphics and Image Processing* 31(2): 156–77

Treisman, A., and Gelade, G. 1980. A feature-integration theory of attention, *Cognitive Psychology* 12: 97–136

Treisman, A., Vieira, A., and Hayes, A. 1992. Automaticity and preattentive processing, *American Journal of Psychology* 105(2): 341–62

Tulving, E. 1972. Episodic and semantic memory, in E. Tulving and W. Donaldson (eds.), *Organization of Memory*. New York: Academic Press: 382–402

Unsworth, N., and Engle, R. W. 2007. On the division of short-term and working memory: An examination of simple and complex span and their relation to higher order abilities, *Psychological Bulletin* 133(6): 1038–66

Velichkovsky, B. M. 1995. Communicating attention: Gaze position transfer in cooperative problem solving, *Pragmatics and Cognition* 3(2): 199–222

Vemuri, S., and Bender, W. 2004. Next-generation personal memory aids, *BT Technology Journal* 22(4): 125–38

Ware, C. 2000. *Information Visualization: Perception for Design*. San Francisco: Morgan Kaufmann

Welford, A. T. 1952. The 'psychological refractory period' and the timing of high-speed performance: A review and a theory, *British Journal of Psychology* 43: 2–19

Wickens, C. D. 2002. Multiple resources and performance prediction, *Theoretical Issues in Ergonomics Science* 3(2): 159–77

Wilson, B. A., Evans, J. J., Emslie, H., Alderman, N., and Burgess, P. 1998. The development of an ecologically valid test for assessing patients with a dysexecutive syndrome, *Neuropsychological Rehabilitation* 8(3): 213–28

Winkler, I. N., Czigler, I. N., Sussman, E., Horváth, J. N., and Balázs, L. S. 2005. Preattentive binding of auditory and visual stimulus features, *Journal of Cognitive Neuroscience* 17(2): 320–39

Wixted, J. T. 2004. The psychology and neuroscience of forgetting, *Annual Review of Psychology* 55(1): 235–69

Wolfe, J. M. 2001. The level of attention: Mediating between the stimulus and perception, in L. Harris and M. Jenkin (eds.), *Levels of Perception: A Festschrift for Ian Howard*. New York: Springer Verlag; draft of 8/2/01, 12:06 p.m. retrieved 15 February 2005 from http://search.bwh.harvard.edu/pdf/Festschrift.pdf

Woods, D. D. 1995. The alarm problem and directed attention in dynamic fault management, *Ergonomics* 38(11): 2371–93

Wright, R. D. 1998. *Visual Attention*. New York: Oxford University Press

Yantis, S. 2000. Goal-directed and stimulus-driven determinants of attentional control, *Attention and Performance* 18: 73–103

Zijlstra, F. R. H., Roe, R. A., Leonova, A. B., and Krediet, I. 1999. Temporal factors in mental work: Effects of interrupted activities, *Journal of Occupational and Organizational Psychology* 72: 163–85

Zimmermann, T. D., and Meier, B. 2010. The effect of implementation intentions on prospective memory performance across the lifespan, *Applied Cognitive Psychology* 24(5): 645–58