

A Survey of Attention Management Systems in Ubiquitous Computing Environments

CHRISTOPH ANDERSON, University of Kassel, Germany

ISABEL HÜBENER, University of Kassel, Germany

ANN-KATHRIN SEIPP, University of Kassel, Germany

SANDRA OHLY, University of Kassel, Germany

KLAUS DAVID, University of Kassel, Germany

VELJKO PEJOVIC, University of Ljubljana, Slovenia

Today's information and communication devices provide always-on connectivity, instant access to an endless repository of information, and represent the most direct point of contact to almost any person in the world. Despite these advantages, devices such as smartphones or personal computers lead to the phenomenon of attention fragmentation, continuously interrupting individuals' activities and tasks with notifications. Attention management systems aim to provide an active support in such scenarios, managing interruptions, for example, by postponing notifications to opportune moments for information delivery. In this article, we review attention management system research with a particular focus on ubiquitous computing environments. We first examine cognitive theories of attention and extract guidelines for practical attention management systems. Mathematical models of human attention are at the core of these systems, and in this article, we review sensing and machine learning techniques that make such models possible. We then discuss design challenges towards the implementation of such systems, and finally, we investigate future directions in this area, paving the way for new approaches and systems supporting users in their attention management.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Human-centered computing** → **Human computer interaction (HCI)**;

Additional Key Words and Phrases: Ubiquitous Computing, Attention Management, Interruption Management, Cognition

ACM Reference Format:

Christoph Anderson, Isabel Hübener, Ann-Kathrin Seipp, Sandra Ohly, Klaus David, and Veljko Pejovic. 0000. A Survey of Attention Management Systems in Ubiquitous Computing Environments. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 0, 0, Article 00 (0000), 27 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

The emergence of computing and communication devices has radically changed the way we communicate and exchange information. It has been estimated that by 2025, the number of unique mobile subscribers will reach 5.9

The authors would like to thank the participants and the organizers of Dagstuhl Seminar 17161 on Ambient Notification Environments for fruitful discussions that helped to shape our research vision. The authors would also like to thank the anonymous reviewers for their helpful and constructive comments that greatly helped to improve the final version of the paper. This work has been co-funded by the Social Link Project within the LOEWE Program of Excellence in Research, Hessen, Germany.

Authors' addresses: C. Anderson, I. Hübener, and K. David, Chair for Communication Technology, University of Kassel, Wilhelmshöher Allee 73, 34121, Kassel, Germany; email: comtec@uni-kassel.de; Ann-Kathrin Seipp, and Sandra Ohly, Chair of Business Psychology, University of Kassel, Pfannkuchstraße 1, 34121, Kassel, Germany; email: ohly@uni-kassel.de; V. Pejovic, Faculty of Computer and Information Science, University of Ljubljana, Vecna pot 113, Ljubljana 1000, Slovenia; email: Veljko.Pejovic@fri.uni-lj.si.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 0000 Copyright held by the owner/author(s).

2474-9567/0000/0-ART00

<https://doi.org/0000001.0000001>

billion, representing over two-third of the world's population. Furthermore, the increased penetration rate of the mobile Internet – over 5 billion estimated by 2025 – in combination with the technological advancements of mobile devices provides information access anywhere and anytime [10]. Such devices come with applications that help us navigate in unknown areas, keep distant social contacts alive, schedule our time, do shopping on the move, and in general help us with tasks that, in the past, were impossible or cumbersome to achieve.

Mobile computing and connectivity, in addition to anywhere-anytime access, provide a means for proactive information delivery. For example, online social networking applications can keep us up-to-date with other users' reactions to our posts; colleagues can inform us about work projects; retailers can send us discount offers as we are passing by shops. This kind of interaction provides seamless support for multitasking and better time utilization. Through mobile notifications, we can receive actionable information *in situ*. For example, an alert about a phone's battery status reminds us to charge it before leaving home; information about a sudden traffic jam down the road helps us adapt our driving route; and a text message stating that a person we are scheduled to meet with will be 20 minutes late allows us to use the extra time more efficiently. Nevertheless, such an opportunity to guide a user towards efficient multitasking often results in fragmenting a user's attention, reducing work performance [60], increasing task error rates [11], inducing stress [64], and even facilitating uninstalls of applications due to annoyed users [86]. With the increasing number of applications and devices we use, the chances for irrelevant and unwanted interruptions rise. Consequently, an average mobile user is faced with about 100 mobile notifications per day, many of which turn out to be irrelevant or distracting [69]. To ameliorate this issue, users may even completely disable all notifications, essentially rendering the whole mechanism designed to support multitasking irrelevant. The decision on when to interrupt a user is left to application developers, and therefore, designing and developing intelligent attention management systems, is crucial to ensure that ubiquitous computing naturally supports, rather than interferes with, our everyday life. Despite the fact that today's mobile applications take a rather *ad hoc* approach to sending notifications, attention management has gained significant research traction among the ubiquitous computing community. A number of recent survey articles investigated attention management systems indicating the rising importance of this field. The existing reviews incorporated concepts from cognitive sciences [96], or focused on technical realizations of anticipatory mobile computing, including interruptions as part of personalized interactions [83]. The survey of interruptibility prediction in ubiquitous systems by Turner et al. is the one most related to our work. Yet it focuses on a meta-analysis of selected interruptibility management approaches and examines specific aspects of these, namely scenario selection (e.g., lab vs. in-the-wild), data collection, and prediction of interruptibility through machine learning [111]. In this article, we rather expand the investigation of attention management systems, following a *combined* approach by providing (i) a holistic investigation of the theoretical underpinnings of human attention, (ii) an overview of the existing approaches to attention management in ubiquitous computing, (iii) guidelines towards designing attention management systems, (iv) recommendations for future research in this field, so to close the gap between the cognitive theory and ubiquitous computing practice.

This article is organized as follows: In Section 2, we first examine definitions of the term *attention* and explore factors that capture it. We further discuss how interruptions in ubiquitous computing environments are created and handled. In Section 3, we present recent theories on how the human brain processes interruptions. Our focus is to bring experiences from cognitive psychology closer to ubiquitous computing developers and is driven by the observation that attention management systems need to *sense* and *understand* the user to enable adaptive attention management. Equipped with sensors that provide a view into a user's contextual environment, and to a certain extent, even into an individual's internal state (e.g., stress [61], emotions [93]), modern mobile devices are well suited for this purpose. Section 4 provides a detailed review of approaches towards sensing and modeling users' behaviors concerning mobile interruptions. We examine different sensing modalities and machine learning approaches used for interruptibility inference, as well as metrics towards evaluating interruptibility models. Finally, in Section 5 we examine the implications that different approaches of interruption management can have

on our health and well-being. Further, we identify discrepancies between the rules that govern our cognitive processes and the current state-of-the-art systems for attention management in ubiquitous computing, such as to interrupt at task boundaries, allow users to rehearse tasks, or provide hints to help users return to the original task after an interruption, setting guidelines for future research in this area.

Key takeaway points

- ★ Attention is captured and steered by external and internal stimuli, while stimulus properties, such as its duration, location, intensity, etc., impact the user's reaction to the stimulus. Multimodal alert types, e.g., sound, light, vibration, and multiple device environments call for a coordinated judicious use of alerting in ubiquitous computing.
- ★ Limited cognitive capacities and threaded task processing imply that, in order to minimize disruptions, interruptions need to arrive at task boundaries or during routine tasks, should allow for task state rehearsal, and support context retrieval through hints presented to the user after an interruption.
- ★ Sensor data from ubiquitous computing devices reveals a user's location, physical activity, collocation with other people, and other information of a user's context that can be related to interruptibility. Next generation wearable devices and personalized machine learning models promise to bring us closer to direct inference of a user's cognitive processes.

2 ATTENTION AND INTERRUPTION: DEFINITIONS AND STRATEGIES

In this section, we review definitions of the terms *attention* and *interruption*. We examine the connection between attention shifting and interruptions and discuss how interruptions can be handled through attention management systems in ubiquitous environments.

2.1 What is Attention?

There is no common understanding of *attention* in the literature. Attention is often considered as a selective processing of incoming sensory information [33], with limited capacity [23] and reactive and deliberate processes [92]. Attention is also referred to as the ability to ignore irrelevant information [24]. The process of selecting stimuli can be voluntary or be steered by external events. The former type refers to goal-driven attention [118], whereas the latter relates to stimulus-driven attention [26]. In this article, we focus on stimulus-driven attention further referring to the definition of Ashcraft et al. who define attention as:

"the mental process of concentrating effort on a stimulus or mental event." [9]

Guided by these definitions, attention is considered as an internal cognitive process that allows an individual to select tasks or information that will be actively processed.

2.2 External vs. Internal Stimuli

Maintaining attention, for example, on a task is often difficult as we are repeatedly confronted with numerous stimuli that compete for our attention. A stimulus, once it captures our interest, leads to an attentional shift changing the selection from a previous stimulus to another. Attentional shifts can be induced by internal or external stimuli. For instance, motivation, thoughts, or emotions represent internal stimuli, e.g., a researcher's desire to work on a scientific paper. External stimuli are related to events originated from the surrounding environment, such as a ringing smartphone drawing our attention from working on a project proposal. Attentional shifts caused by internal stimuli are considered as internal interruptions, whereas external stimuli lead to external interruptions [30, 72]. This paper focuses on automated systems that manage external interruptions (e.g., notifications) supporting individuals in maintaining their attention on tasks and activities. In the following section, we describe the structure of interruptions and their effects on individuals, activities, and tasks.

2.3 Interruptions and Distractions: Structure and Handling

Interruptions are investigated in various research fields, e.g., the psychology of attention, or Human-Computer-Interaction (HCI). In fact, the term interruption is defined differently, according to the requirements of the respective research domain¹. A particular definition in the field of psychology states that an interruption is

"an introduction to a new task or tasks on top of the ongoing activity, often unexpectedly, resulting in conflicts and loss of attention on the current activity, failing to resume the work where it was interrupted." [72]

This definition implies that an interruption affects an individual's current activity, resulting in the loss of attention and the inability to resume the activity where it was interrupted. Interrupting a task means to abandon a current task before finalization and diverting the attention to a new or different task [30]. The definition of interruption is closely related to distraction, which is researched in the field of road safety (see [95] for a review). The difference between distractions and interruptions is that the former are encountered stimuli intended to be ignored [25] and filtered out by a top-down suppression of signals from the prefrontal cortex [20, 21]. The latter are stimuli that represent aspects of a secondary task, resulting in a reallocation of cognitive resources [25]. For the scope of this article, we refer to the term *interruption*, further investigating the structure of interruptions as well as handling strategies and approaches.

Structure of Interruptions & Handling Strategies. Interruptions are not single events that interfere with an individual's current task or activity. Instead, they are embedded in a complex process that incorporates phases before, during, and after an interruption (see Figure 1). Before an interruption, an individual is focused on a task, known as the pre-interruption phase [18], or primary task performance [4]. At a certain moment (end of Step *i* in Figure 1) an individual perceives an interruption (e.g., phone call). This realization is followed by an *interruption lag*, a brief transitional period that precedes a pending interruption. In this period, an individual is aware but has not yet engaged the interruption, providing an opportunity to complete thoughts or to negotiate trending activities [3, 4]. For example, knowing that a colleague has already entered the office, a researcher might finish a sentence within a research paper first, rather than responding to the interruption immediately. After perceiving an interruption, an individual decides whether the interruption should be handled momentarily, later or not at all. Four general strategies for dealing with interruptions have been proposed in the literature [18, 65, 66]. In particular, an individual might decide to stop working on a task to focus on the interruption *immediately*. Also, an individual can *negotiate* an interruption, for example, by first finishing a sentence of a research paper before responding to a colleague. Interruptions can also be *scheduled* or *mediated*. In the former case, an interruption is handled at a specific time or within a given period (e.g., every 15 minutes). In the latter case, an interruption is mediated by another person, proxy or system that interrupts an individual indirectly and decides *when* and *how* to interrupt. Before resuming, there is a lag between the interruption and the first action on the primary task. The *resumption lag* is a time span, where an individual rather needs to collect thoughts or hints about the last known status of the primary task before resuming to it [4], even when an interruption was only perceived and not actively handled. These procedures are undertaken in the post-interruption phase, specifically in *step r* depicted in Figure 1.

2.4 General Effects of Interruptions

Most of the previous research agrees that interruptions result in prolonged completion times and increased error rates for *primary tasks* that participants initially worked on [11]. For instance, Speier et al. analyzed the effect of interruptions in a laboratory study using computer-based tasks and concluded that both, the error rate and the completion time of the primary task, are negatively affected by an interruption, but only if the interrupting

¹We refer the reader to [65] for a review.

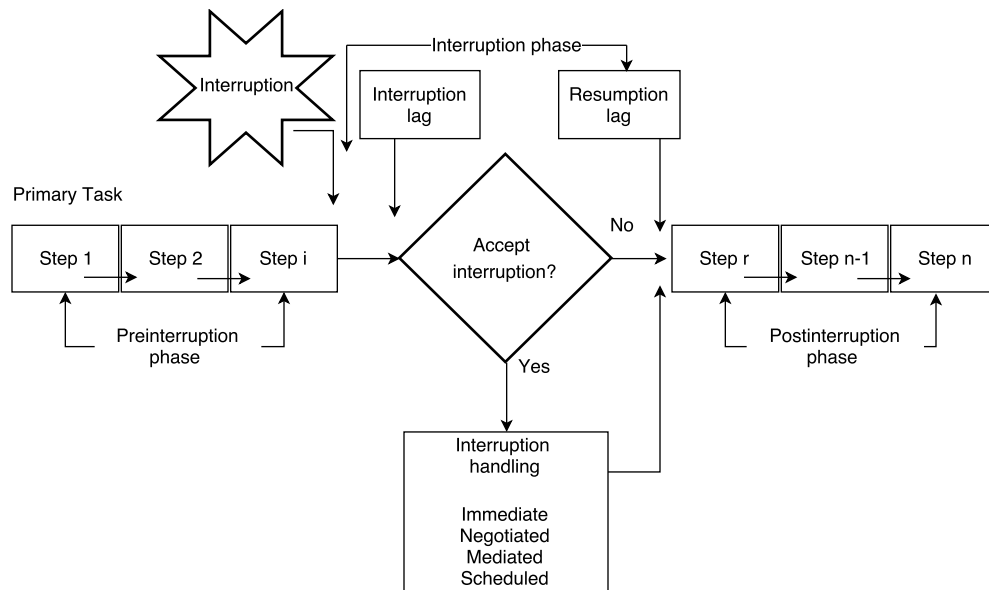


Fig. 1. Interruption model adapted from [18]
©2007 Wolters Kluwer Health | Lippincott Williams & Wilkins

task is complex enough [105]. Cutrell et al. state that the completion delay originates from switching from the primary task to the one signaled by the notification, as well as from switching back to the primary task [27].

Regarding the effect on the secondary task, findings are less univocal. Bailey and Konstan's experiments show that secondary tasks are performed faster when delivered as an interruption to the primary task, rather than stand-alone tasks [11]. The authors deduce that this comes from a user's increased motivation to get back to and complete interrupted primary tasks as quickly as possible. At the same time, in cognitive load research, the reaction time, accuracy, and error rate on the secondary task are assumed to deteriorate, not improve, and the secondary task performance is often used as a measure of cognitive load imposed by the primary task [81]. The subtle difference might arise from the way experiments are designed: interruptibility research often assumes that users are task-switching so that cognitive resources are focused on a single task at all times, while cognitive load research assumes that tasks are done in parallel so that resources are shared among multiple tasks.

2.5 Interruptions in Ubiquitous Computing

The evolution towards ubiquitous computing has drastically changed the way we communicate and receive notification. Computing devices are now ubiquitous, personalized and provide always-on connectivity. Devices such as smartphones allow their owners to operate in parallel spheres by, for example, being in a phone call with one person, instant messaging with another while performing an unrelated task in the physical environment. Although, as early as in 1991 Mark Weiser started calling for "*stealth*" ubiquitous computing, where devices quietly blend with the user's lifestyle [116], as of today, we are still far from a realization of this vision. The rapid interweaving of ubiquitous computing in our everyday life facilitate numerous channels for communication while at the same time leaving only a few barriers for attention capturing. The net effect of an uncontrolled exposure to interruptions is negative, especially in terms of task completion [11], performance [60], and emotional states [64].

Whereas the effect of interruptions on task performance was the primary focus of early research on interruptibility, understanding the impact mobile notifications have on a user's internal state emerged as a critical research question in ubiquitous computing. Increased stress and frustration caused by computer-based interruptions have been observed by Adamczyk and Bailey [1]. The authors found that the timing of interruptions within task execution has a significant effect on an individual's perceived amount of stress and frustration. More stress and frustration is induced when interruptions are timed between the execution of interrelating subtasks and less if timed after subtask completion². A study by Mark et al. confirms stress-inducing effects of instant messaging and phone interruptions [64]. Furthermore, Kushlev et al. investigate effects of interruptions caused by notifications on a group of 221 participants who are exposed to the maximum level of mobile notifications for one week and for another week assigned to the minimum amount of such notifications [58]. Their findings show a positive correlation between interruptions and symptoms of hyperactivity and inattention of individuals.

With ubiquitous connectivity, there is a need to be up-to-date with the events of interest. Internal interruptions (see Section 2.2) are increasingly common and new behavior patterns, such as constant notification checking, have appeared. Kushlev and Dunn demonstrate the negative effect of frequent email checking on users' well-being [57]. In another study about participants, who deliberately disabled all push notifications on their smartphones, Pielot and Rello found that the lack of notifications indeed increases participants' self-reported productivity, but at the same time leaves the participants anxious about not being responsive as expected and being less connected with their social contacts [90]. Finally, an earlier study of instant messaging users finds that without these messages users fallback to in-person interruptions which in certain work environments lead to even more stress [38]. These examples imply that ubiquitous computing does not present an inherently negative technological determinant. However, attention management systems need to be designed carefully to facilitate multitasking, minimize anxiety and the sense of social isolation, while ensuring that users are not overloaded with information and distracted in their work.

3 MULTITASKING: HOW INTERRUPTIONS ARE PROCESSED

A major issue with attention management in ubiquitous computing environments is the dissonance between interruptions, often signaled by mobile notifications, and the flow of cognitive engagement of users. The human brain, although sophisticated, is still limited when it comes to processing concurrent signals and tasks. In this section, we elaborate on the theoretical frameworks that explain how our brains handle tasks competing for the limited cognitive resources we have. We start with the unified theory of multitasking developed by Salvucci and Taatgen [97], which is itself based on Anderson's Adaptive Control of Thought-Rational (ACT-R) architecture [7]. This experimentally validated framework provides an analytic model of the human cognitive capacities and their use. We then examine how cognitive resources are used when individuals face concrete tasks and turn to the cognitive load theory to explain how interruptions interplay with ongoing tasks of varying complexity. The result of the analysis is the identification of the theoretically most suitable moments to interrupt individuals. Finally, we consider extensions to the presented models that indicate the importance of other factors, such as the content of the interrupting task and the social pressure on the user, or the way interruptions influence a user's reaction. These findings are further discussed in Section 5.3, where we derive practical guidelines for designing attention management systems that are based on theoretical findings.

3.1 Theory of Multitasking

Fueled by the advances in information and communication technologies, the number of information and task sources keeps increasing, straining the limited cognitive resources we have. The understanding of cognitive processes is essential for efficient attention management in ubiquitous computing environments. In this article,

²In Section 3, we discuss the cognitive background of this finding in detail.

we distill from the vast research in psychology and neurology that is concerned with multitasking and present the theoretical foundations of concurrent task processing.

3.1.1 Threaded Cognition. Threaded cognition is one of the most recent theories that explains multitasking via parallel threads competing for the same processing resources [97]. The concept is quite similar to the way multithreading is implemented in modern single-CPU/core computers:

- Each active thread is associated with its own goal.
- Resources execute processes exclusively in service of one task thread at a time.
- Threads acquire and release resources in a greedy, polite manner.
- When multiple threads contend for the procedural resource, the thread with the highest urgency proceeds.

The resources these threads compete for are cognitive resources, including declarative resources representing static knowledge that can be recalled (or forgotten), procedural resources representing procedural skills as goal-directed production rules, and perceptual and motor resources, allowing information acquisition from the environment and actions in the environment, respectively. Even the simplest tasks usually require more than a single kind of a resource, and complex patterns of interference may arise when two or more tasks are to be executed in parallel. Take this simple example from Borst et al.'s experiment [16], depicted in Figure 2. One task requires the user to press a key when a visual stimulus is shown (white boxes), while the other requires a vocal response from the user after an auditory stimulus (gray boxes). The x-axis represents time, and the boxes correspond to the period in which a resource is used (either by one task or the other). The tasks commence by activating a production rule related to a corresponding stimulus and continue with the encoding in either the visual or the aural module. We see that the interference (marked with **A**) caused by the contention for the procedural module prevents the perfect in-parallel execution of the task. Such interference can lead to delayed completion times and errors.

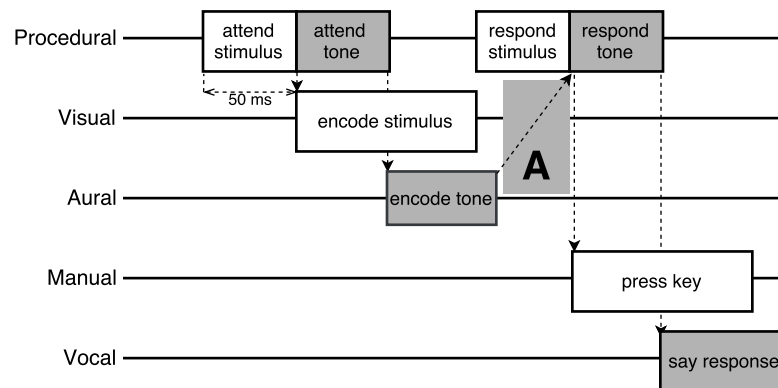


Fig. 2. Processing stream in threaded cognition adapted from [16]

3.1.2 Computational Modeling of Multitasking. How exactly the brain harnesses cognitive resources, how much processing time each operation takes, and how the resources map to different parts of the brain is explained by cognitive architectures, for instance, the ACT-R architecture developed by Anderson [7]. This theory describes the procedural memory regarding production rules, while declarative knowledge gets retrieved in chunks as necessary. Different modules, such as visual, aural, and manual, represent interfaces both for perceiving the physical world, as well as acting in it. However, the central production system is not sensitive to everything that

happens within these modules. Instead, it can only respond to a limited amount of information that is deposited in the buffers of these modules. In other words, the ACT-R architecture puts a limit on human attention. Similarly, only the currently retrieved facts, and not all the information in our long-term memory, are readily available for the production system. In addition to explaining cognitive processes on the conceptual level, ACT-R also provides a computational tool, enabling fine-grain simulation of task duration³. This theory has been verified through a number of studies with different task types, from mathematical problems to driving simulations. Salvucci and Taatgen use ACT-R for the implementation of their threaded cognition theory and extend the computational model to support multitasking [97]. They confirm that the postulates of the threaded cognition theory realistically describe cognitive processes that happen when humans handle concurrent tasks.

To understand how interruptions affect cognitive processes, we have to examine concepts associated with the primary and the interrupting task. According to the ACT-R theory, when recovering from an interruption, primary task chunks that were stored in declarative memory have to be retrieved. The easiness at which they are retrieved is related to their (i) base-level activation, a measure of a chunk's inherent likelihood of being recalled, and (ii) strength of association, a measure of how the current environment context facilitates cueing of the particular chunk. For attention management system designers, this presents a clear case for including hints about the interrupted task to facilitate task resumption.

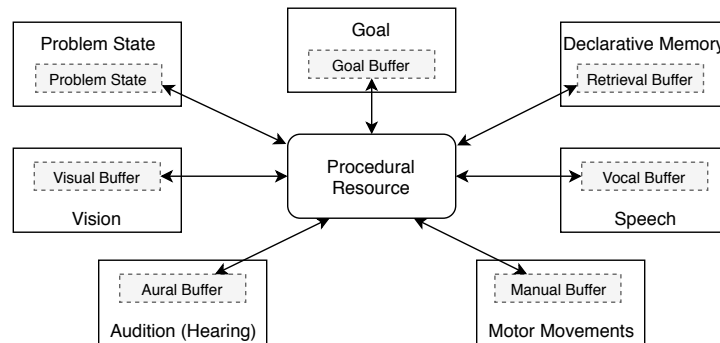


Fig. 3. ACT-R Architecture of cognitive resources adapted from [98]
©2011 Dario Salvucci & Niels Taatgen

Following the ACT-R theory, each task has an associated goal that can also include subgoals. This is a cornerstone of Altmann et al.'s theory called Memory-for-Goals [3], according to which a task interruption results in the current active goal being stored in declarative memory. Task suspension and resumption processes can then be broken down into the core processes of rehearsal (or strengthening) and recall. When the task is suspended, the stored associated task goal decays, albeit its activation is raised with rehearsal. According to Altmann et al., the cognitive system must strike a balance between rehearsing the suspended task so that it can be easily retrieved when needed, but at the same time not activate it to the point where it interferes with the newly-attended task. The authors suggest that goals are preferably retrieved when they are associated with a cue, and when the goal-cue relationship is built just before the goal was suspended.

The Memory-for-Goal theory was also used to explain the difference in the time needed to revert to the primary task after an interruption depending on the interrupting task duration and complexity. In the experiment by Monk et al., users performed a complex Video-Cassette-Recorder programming task, while getting interrupted with tasks of different complexity [73]. Resumption lags were longer for complex tasks than for simple tasks.

³An open-source Lisp implementation of the framework is available at the following URL: <http://act-r.psy.cmu.edu/software/>

Also, longer interruptions resulted in longer task resumption. This can be explained with the time decay of goal activation proposed in [3]. Monk et al. discuss immediate design guidelines grounded based on these results, such as an in-car navigation system that should strive to minimize the interruption complexity, not necessarily the interruption duration, as simple interruptions lead to relatively fast primary task resumption, irrespective of the duration. However, neither of the theories investigate nuances in interruption complexities or why the resumption lag does not seem to scale to a certain level with interruption complexity.

3.1.3 Memory for Problem State. The explanation of why complex tasks cause stronger interferences is provided by Borst et al. [17]. The authors build upon the threaded cognition theory. Threaded cognition allows for multiple parallel goals, and thus multiple tasks (threads), to be active. This translates into the assumption that the goal module in ACT-R can represent several goals at the same time. Depending on its complexity, a task may or may not require *the problem state*. The problem state resource is used to maintain intermediate mental representations that are necessary for performing a task. Borst et al. give an example of an algebra problem, such as $2x - 5 = 8$, where the problem state can be used to store the intermediate solution, i.e., $2x = 13$. There is also physiological evidence of the existence of the problem state, as neuroimaging experiments find that the transformation of mental representations correlates with the blood oxygen level-dependent activity in the posterior parietal cortex [8, 104]. At the moment of interruption, if both the primary and the interrupting task are complex, the problem state of the primary task is stored in declarative memory where it starts to decay. On the return from the interruption, the primary task's problem state has to be retrieved from the memory, yet, this process is tied with delay and errors. Even more delays and errors are to be expected if the problem state has lingered in the declarative memory for a longer period of time. Experimental results with tasks of differing complexities confirm this theory and show that when either the primary or the interrupting task is simple enough not to require the problem state, the resumption time and the task errors are minimized.

3.2 Beyond Interruption Timing

The theories outlined in this section are bound to have substantial implications on how attention management systems might evolve (see Section 5.3). However, these theories do not present a comprehensive explanation of interruptibility. This is particularly true for attention management in ubiquitous computing environments where other aspects, such as the type of the communication channel, the content of the interruption/message, the frequency of interruption, the relationship between the sender and the receiver, the alert type, and the social environment, may play a deciding role in how an interruption is handled.

Considering the *interruption content*, Speier et al. show that interruptions containing information dissimilar from the primary task take longer to complete than those with content related to the primary task [105]. Since the success of the retrieval from declarative memory depends on the context, the closer the sought for information is to the currently active one, the easier it is to retrieve it. Addas and Pinsonneault consider interruptions that are either relevant to the task at hand, or not, and that are either actionable or just informative [2]. An actionable irrelevant interruption leads to task errors and longer completion times. Relevant information, on the other hand, helps with the task quality but may affect the time to finish the task. The findings are important for project management, for example. If one's task responsibilities are defined too narrowly, many interruptions that are not directly related to the task will be perceived as interruptions. Expanding the task boundaries allows for a more fluid definition of relevant interruptions.

Grandhi and Jones propose a theoretical framework that explains interruption handling decisions, and takes into account the user's *cognitive, social* (the place the interruptee is in, the social aspect of the place, and people present at that place) and *relational context* (who is the message from, what is it about, and the history of the interruptee-interrupter relationship) [42]. Through experimental validations, the authors show the multifaceted nature of interruption handling. Even though the effects of interruptions on task performance or appropriateness

in a social setting may be negative, individuals may still decide to handle an interruption. The authors argue that people may consider not just how an interruption affects one's local task, but also how it may affect other tasks and interpersonal relations. Consequently, the authors claim that instead of an automated interruption management system, efforts should be turned to User-Interface (UI) designs that convey the necessary relational context to the interruptee and facilitate one's interruption handling decision process.

Personal devices such as smartphones or personal computers allow experimental validation and further empirical investigation of factors related to interruptibility. A large-scale study of mobile notification usage by Shirazi et al. shows that users assign different importance to notifications triggered by different application categories [103]. This is unlikely to be related to the specifics of the application's notification management but is probably conditioned on the application design and the social role that a user assumes in different applications (e.g., Skype for business contacts, Whatsapp for family chats). That mobile computing users indeed handle interruptions from different contact types (e.g., family, friends, work) differently has been shown in a study by Mehrotra et al. [67]. Pielot et al. found that users opt for immediate handling of potentially disruptive interruptions just to conform to social expectations, demonstrating the social pressure in notification handling [87].

Finally, physical properties of stimuli determine whether and how the attention is going to be steered. Two independently conducted studies have confirmed the difference in interruption handling times depending on whether a notification vibrates, blinks or sounds an alert [70, 87]

4 ATTENTION MANAGEMENT SYSTEMS

In this section, we examine functional principles of attention management systems, in particular, systems that (partly) seize the theories and frameworks mentioned before. According to Bailey and Konstan, an attention management system is defined as

"a system that computationally seeks to balance a user's need for minimal disruption and the application's need to efficiently deliver information." [11]

In other terms, an attention management system might aim to steer an individual's attention in case of an interruption by choosing a suitable output modality gently or even postponing interruptions to a time where it could be handled by a user. For example, if we consider push notifications as interruptions, this means that a system automatically decides *when* and in *which ways* an individual is notified. How does a system differentiate between disruptive and non-disruptive interruptions? The answer lies in a computational model that holds characteristic information about potential disruptive interruptions. In the following, we examine implementations of underlying interruptibility models as well as designs of practical attention management systems.

4.1 Designing Attention Management Systems

Sensing, actuating, and machine learning are at the core of attention management systems. These concepts interplay through different processing stages that connect raw data from the physical environment, over interruptibility modeling to attention management. These stages include *sensing*, *processing*, *inferring*, *modeling*, and finally *managing* interruptibility as illustrated in Figure 4. To demonstrate key challenges and implementation issues, take a running example of an attention management system inspired by Attelia II [77], depicted in Figure 4. Attelia II monitors acceleration signals and application usage to infer posture and locomotive activities of users, e.g., walking, sitting, and standing as well as a user's device usage. The system builds an interruptibility model that is based on theoretical postulates (see Section 5.3) assuming that times between tasks represent opportune moments for interruptions. Finally, the system detects natural breakpoints [74] in physical and user-interface-based activities and postpones notifications to these moments.

4.1.1 Sensors and Features. Equipped with microphones, light, acceleration, temperature, and other sensors, as well as communications interfaces, mobile computing devices can *perceive* their surrounding environment

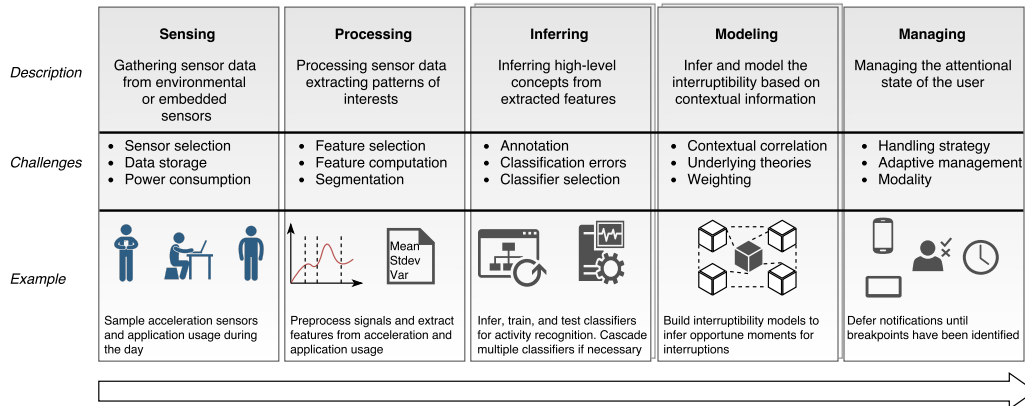


Fig. 4. Stages of an attention management system including sensing, processing, inferring, modeling, and managing. Key challenges along with an example are depicted for each stage.

with the help of machine learning models. The key premise for building meaningful models is to select features, i.e., sampled data representations, that correlate with the classes of interest [32]. In Figure 5, we sketch the process of deriving features from sensed data. Raw sensor data (e.g., acceleration) is recorded, stored and then preprocessed, for example, by applying segmentation, sanitization, and normalization to remove artifacts such as outliers, noise or missing data. Features that correlate with targeted classes are then computed on preprocessed data. In this regard, a selection of the most informative features for a particular scenario is required given the vast amount of possible features. In general, such a selection is based on features that have been found useful to discriminate patterns in the domain at hand⁴. Table 1 provides an excerpt of commonly used features in attention management systems. The sensors and features listed here are not exhaustive and are intended to be a starting point for designers of attention management systems.

Besides sensors embedded in mobile devices, attention management systems may also draw information from personal computers as well as from sophisticated psychophysiological sensors. Psychophysiological sensors are used to extract information correlating with an individual’s cognitive state (e.g., mental workload), whereas mobile devices and personal computers are found to provide additional characteristics about an individual’s interruptibility, for example, related to activities, or application usage. The process of extracting features from sensed data is exemplified by Attelia II which uses acceleration sensors to infer an individual’s posture and locomotive activities [77]. In the first step, raw acceleration signals are recorded and stored on the device. The system then computes 22 time and frequency-domain features – e.g., acceleration mean, variance, or correlation – frequently used in the field of activity recognition that corresponds with an individual’s movement.

4.1.2 Contextual Information and Machine Learning Techniques. The major theoretical determinants of interruptibility, such as the interplay of a stimulus signaling an interruption, a user’s current task engagement, and the complexity of the interrupting and the interrupted tasks, were investigated in Section 2.3. As attention management systems have often developed independently from underlying theories, they use contextual descriptors and features that reflect, *but not necessarily determine*, the interruptibility state of an individual. In an *optional* stage of inference, machine learning models use features to recognize contextual information of users and their environments, such as physical activities, semantic locations, or even emotions. Figure 6 illustrates this process - features are first further pruned and compressed, for example, by applying the Principal Component

⁴We refer the reader to Section 4.3 of [83], which summarizes the most commonly used features in different inference domains.




Fig. 5. Sensing and feature computation: From sensors to features

Analysis [117] or information gain [44] to select a set of meaningful features. The resulting feature set is used to build machine learning models. As contextual information can be of different levels of abstraction, a plethora of contexts can be extracted through machine learning amplifying the possibilities of interruptibility modeling. A key challenge in the field of attention management, therefore, is to identify contextual descriptors that correlate with an individual's interruptibility, craft machine learning models that can infer these contexts, and finally tie them into an interruptibility model.




Fig. 6. From features to contexts: Deriving contextual information from sensor data.

Commonly used contextual information that correlate with an individual's interruptibility cover various aspects of life.⁵ It is worth to note that not only physical activities [46, 77, 99] or interactions [35, 91], but also more expressive concepts such as complex activities [54], engagement levels [84], even personal traits [64, 119] are frequently used in attention management systems. To infer such contextual information, various machine learning algorithms e.g., J48 [77, 99], K-Means [54], or Support Vector Machines [99] are typically trained and tested⁶. Given the vast amount of classifiers, a selection of classifiers depends on the amount and type of training data, on the processing power of the device, and above all on the classification performance. Statistical tests, e.g., Wilcoxon signed ranks or Friedmann tests, can help to evaluate and compare classification performances [31].

4.1.3 Interruptibility Models and Proxies. How to differentiate between disruptive and non-disruptive interruptions? The answer to this question lies in computational models which hold hypotheses about an individual's interruptibility. These models are either based on concepts and theoretical frameworks that have been constructed to explain the nature of human interruptibility (see Section 3) or purely data-driven, based on the results of mining

⁵We refer an interested reader to the extended interruption taxonomy proposed in [39, 108].

⁶We refer the reader to [15] for an introduction to pattern recognition and machine learning.

Table 1. Sensors and features commonly used in attention management systems to infer interruptibility.

Domain	Sensors	Features
<i>Physical Activities</i>	Acceleration [46, 53–55, 77, 99, 114]	mean, energy, entropy, correlation, variance, ratio {minimum, maximum}
	Orientation [91, 112]	mean pitch angle, mean roll angle, flat, upright
	Gyroscope [114]	mean crossing rate
<i>Device Interaction</i>	Microphone [14, 35, 40, 54, 62, 86]	mean, variance of fourier coefficients, signal power, zero crossings, cepstral coefficients, noise {true, false}, ratio {voice, noise, silence}, spectral {centroid, diffusion}
	Application usage [36, 69, 76, 77, 86, 109, 121]	application name, window changed {state, content}, interface statistics (e.g., min, mean, max), skype status, number of applications per time frame,
	Proximity [89, 91]	screen covered
	Light [40, 70, 112]	intensity, ambient light {dark, dim, light, bright}
	Volume [78, 86, 89, 112]	unknown, silent, vibration, or sound
	Mouse and Keyboard [36, 51, 109, 121]	click events, {moved, scrolled} pixels, keystrokes, most common key events {pressed, released, traversed}
	Calendar [114, 121, 123]	events {free, meeting}
	Battery [40, 112, 113]	charging {true, false}
<i>Location</i>	GPS [35, 40, 69, 99, 119, 119]	longitude, latitude
	WiFi [40, 54, 69]	connectivity, coarse location
<i>Temporal</i>	Time [35, 78, 86, 112, 123]	time of day {morning, afternoon, evening, night}, day of week, weekend {true, false}
	Events [51, 89, 119]	time since: notification {received, viewed}, screen {on, off, covered, uncovered}, time to {react, complete}
<i>Psychophysiological</i>	Electroencephalography (EEG) [37, 43, 122]	brain wave frequency bands (e.g., α (8–12Hz), β (12–30Hz), γ (30–80Hz), δ (0–4Hz), θ (4–8Hz)), variance of θ power, median of β power, {min, max} attention, {min, max} meditation, fractions of wave bands { $\Delta(\alpha/\beta)$, $\Delta(\beta/\gamma)$, $\Delta(\gamma/\theta)$ }
	Electrodermalactivity (EDA) [37, 43, 94, 122]	{min, max} peak amplitude, Δ number of phasic peaks/min, Δ mean phasic peak amplitude, Δ mean skin conductance level, skewness, {mean, variance, median} galvanic skin response
	Electrocardiograph (ECG) [94, 99]	{ratio low/high-frequency components} of heart beat and heart beat frequency, mean, median, stdev, percentiles {10th, 25th, 75th, 90th}
	Photoplethysmograph [122]	{mean, sum, max,} peak amplitude BVP, mean heart rate, percentage difference of interbeat interval {20ms, 50ms}
	Eye tracker [37, 43, 94]	{mean, median, stdev,} saccade duration, {mean, variance, median} pupil diameter, number fixations/min, {mean, median, stdev} fixation duration
	Respiration [94, 99]	duration of inhalation/exhalation, mean, median, stdev
	Temperature [43, 94]	{mean, variance, median} heat flux, temp. difference {nose, neck}

traces describing human behavior in different situations. Examples of the former are attention management systems that interrupt at task boundaries (breakpoints) [1, 48, 77] or when task engagement is low [84]. These approaches are supported by underlying theories claiming that interruptions may require an exclusive resource – the problem state – which should not be occupied by the primary task at the time of the interruption [17]. The latter, data-driven approaches, are realized by the abundance of sensor data gathered from mobile and personal computing devices. InterruptMe, for example, uses time-of-day information together with sensor readings from GPS and acceleration to build a model that explains user behavior concerning interruption handling [82]. The

approach is not based in any particular theory, but might reveal links with existing theoretical frameworks after further investigations.




Fig. 7. Cascading classification: Combining multiple classifiers to infer interruptibility.

To construct interruptibility models, features carrying information sensed from the environment, can either be used directly (e.g., InterruptMe [82]), or through further inference of contextual information (e.g., Attelia II [77]). This *cascading classification* – a subsequent classification is trained on the output of a previous one – facilitates different kinds of contexts, in particular, different levels of abstraction that augment interruptibility models (see Figure 7). For example, Attelia II classifies posture and locomotive activities and use this information to train a second classifier that infers breakpoints in physical activities. Combined with a breakpoint detection approach based on application usage, a final classification model is trained to infer interruptibility.

In Table 2, we list a selection of interruptibility models that are implemented in attention management systems. For example, an individual’s engagement in tasks is investigated in [84] to predict interruptibility. The authors argue that individuals skilled in the task that is interrupted, feel less likely disturbed by an interruption. A possible explanation is that skilled individuals need less time to reconstruct the primary task state after an interruption due to stronger links in declarative memory. However, interruptions, during tasks that require a significant amount of concentration, are more likely perceived as disruptions. The reason for this might be that interruptions within periods of high mental workload – induced by inexperience or difficult tasks – result in higher resumption costs [17] forcing an individual to spend more cognitive resources to resume to the primary task. The connection between mental workload and disruptiveness of interruptions has also been shown in [22, 48, 49].

Different metrics and annotation techniques to measure an individual’s interruptibility can be applied to evaluate respective models. For example, explicit measures include *experience sampling methods/ecological momentary assessments* (ESM/EMA) where individuals are asked to rate their current experience *in-situ*. In this regard, the NASA-TLX is frequently used, in particular, to assess models that are based on cognitive theories and frameworks such as mental workload [122], level of task engagement [84], or level of interruptibility [119]. NASA-TLX is a multi-dimensional rating tool which comprises six weighted scales including *mental*, *physical*, and *temporal demand*, as well as *performance*, *effort*, and *frustration* [45]. However, the assessment of subjective workload using questionnaires may not be feasible in ubiquitous computing environments, where multiple tasks are performed simultaneously. Rather than using *post hoc* questionnaires, implicit measures, including types of reaction, or time to react are used to evaluate attention management systems instead [69, 89]. For example, the authors in [69] measure the time taken to respond to a notification. If an individual responds within a period of 10 minutes, the moment is classified as opportune and inopportune otherwise. The selection presented in Table 2 is not meant to be a complete and comprehensive overview, given the fact that ubiquitous environments provide a variety of possibilities for modelling interruptibility. However, this overview points out that building interruptibility models in the field of attention management can benefit from interdisciplinary research, for example, by incorporating findings from psychology or cognitive science (see Section 5).

4.1.4 Managing Attention. Managing interruptions, once identified by the underlying algorithms and models, means that a system follows specific strategies to deal with interruptions (see Section 2.3). In the past years,

Table 2. **Interruptibility models and metrics of attention management systems.**

	Model	Approach	Metric
Okoshi et al. [75–78]	Breakpoints between physical activities/application usage	Perceived mental workload is lowered by deferring interruptions to natural breakpoints in physical activities/application usage	Breakpoints via self-annotated application usage, mental workload via ESM: NASA-TLX [45]. Likelihood of physical breakpoints via ESM: 10-Point Likert scale. Engagement to content via response time, click rate.
Pejovic et al. [82]	Broader contextual information of individuals including activity, location, and engagement.	Interruptibility is reflected by the context in which an individual is at the moment of interruption.	Reaction presence: respond to a notification, time to reaction: respond to a notification within a time period, and sentiment towards notification via ESM: {not at all, a little, some, very much}.
Pejovic et al. [84]	Task engagement	The degree to which an individual is engaged in a task determines interruptibility.	Levels of engagement: {interesting, challenging, concentrated, and important}, measure of engagement via ESM - 4-Point Likert scale: {not at all, a little, somewhat, and very much}. Questions: Is your current activity: {interesting, challenging}? How well are you concentrated on your task? Is the activity important for you?
Pielot et al. [89]	Attentiveness	The degree to which individuals pay attention to incoming notifications reflect interruptibility. Attending means that an individual either opens the application that received a notification or opens the notification drawer.	Attentiveness: {very high, high, low}, measured implicitly by the mobile phone usage.
Sarker et al. [99]	Availability	Interruptibility is determined by an individual's capability of engaging tasks or activities.	Availability via EMA [102], multiple times a day.
Yuan et al. [119]	Reaction & Interruption intensity	Interruptibility is determined by first predicting an individual's reaction and second the intensity of an interruption.	EMA [102] about participants' behaviors, physiological and psychological states, including 42 items, current activity, and mood {pleasant, unpleasant} etc.
Züger et al. [122]	Mental workload	The perceived mental workload determines individuals' interruptibility.	Mental workload via ESM - 5-Point Likert scale: {(1) very-low, (5) very-high}, perceived disturbance of interruptions: {(1) not all disturbing, (5) very disturbing}, interruptibility level: {(1) highly interruptible, (5) not all interruptible} Mental workload measured with psychophysiological sensors (see Table 1).

research focused on attention management systems which mediate and negotiate interruptions. Another strategy that has been used in attention management systems is the concept of *mitigating* interruptions. Approaches following this strategy, for example, change the modality of a disruptive notification (e.g., vibration, silent, visual) rather than postpone them to an appropriate point in time. In Table 3, we list a few examples of applied management strategies in attention management systems. For example, Zulkernain et al. proposed an intelligent mobile interruption management system that utilizes a decision tree to pick an interruption mode (ring, silent, vibrate) based on the user's context automatically [123]. Rather than choosing alert types, other approaches focus on the utilization of Internet of Things (IoT) infrastructures [56, 115]. The authors in [56] propose a system that can track the location of an individual to deliver notifications to nearby output devices such as televisions, personal computers, or even ambient lights. Rather than deferring or changing the representation of interruptions, a few systems indicate the interruptibility status of individuals. For example, the system proposed in [121] indicates the interruptibility of employees via colored lights. Their findings show that employees perceive fewer interruptions over the day as they take the light status into account before interrupting colleagues.

However, an attention management system should adapt its handling strategy to a user's current needs. This means handling strategies, as well as underlying interruptibility models, should not be static and predefined beforehand. In this regard, choosing the *right* strategy remains one of the key challenges as an individual's behavior, and preferences might change over time. Different strategies and an adaptive mechanism that reacts to changes according to an individual's preferences are essential for such systems. One possible solution for adaptive attention management is proposed in [67]. The authors introduce a mechanism that *mines* association rules by analyzing notification preferences of users using notification titles and different contextual information (e.g., location, activity, and time). Based on the user's action (dismiss or accept), the system generates a set of rules that reflect a user's attitude towards new notifications based on the current context. This approach is not only intuitive for the user, who consequently might put higher trust in an attention management system constructed in such a manner, but also interesting for further examination of theoretical frameworks that explain human interruptibility.

Table 3. **Examples of attention management in ubiquitous computing environments.** In addition, approaches following the given management policies are listed.

Strategy	Management Policy	Approach
<i>Mediating</i>	Deferring to breakpoints	Physical activities [1, 76, 77] Device interaction [78, 91] Tasks & Mental Workload [11, 50, 51]
<i>Mitigating</i>	Choosing output device Choosing output modality	Estimating device for receiving notifications [115] Using IoT infrastructure for notifications delivery [56] Inferring preferred output modalities [123]
<i>Indicating</i>	Indicating interruptibility	Showing interruptibility status [13, 59, 121]

4.2 Attention Management Systems Redefined

Attention management encompasses more than balancing an individual's need for less disruptions and efficient information delivery. Attention management systems actively support individuals in managing interruptions and maintaining concentration on tasks and activities by constantly *sensing*, *modeling*, and *managing* interruptibility. Models which help to determine interruptions, partly reflecting the actual state of an individual's interruptibility, represent a key factor of attention management systems. Based on our findings in Section 4, we define attention management systems as

"systems that sense, model, and manage the attentional state of a user. Managing the attentional state is considered as any action of the system, which supports an individual to maintain its concentration on a task or activity."

Key factors of attention management systems

- ★ Sensing, actuating, and machine learning are at the core of ubiquitous attention management systems. Sensing the physical environment allows for feature computation to discriminate concepts and factors of interruptibility.
- ★ Interruptibility models hold assumptions and representations of an individual's interruptibility. Such models are based on features and contextual information that correlate with an individual's state of interruptibility.
- ★ Based on the output of an interruptibility model, attention management systems follow different management strategies. Mediating and mitigating interruptions are the most commonly used strategies.

5 IMPLICATIONS AND PERSPECTIVES

Our attention is one of the most precious resources pervasive applications compete for. To protect individuals from adverse effects of attention steering, we have to ensure that future attention management systems respect

ethical norms and work towards the well-being of their users. For the latter, it is crucial that attention management systems are designed in accordance with human cognitive processes and evaluated in long-term studies.

In this section, we discuss ethical and well-being implications of attention management. We highlight the lack of ethical guidelines when devising attention management systems and the absence of long-term studies on their impact on human well-being, and present suggestions for ameliorating this. We then identify untapped opportunities that arise from our review of cognitive psychology literature and identify key missing features that should be present in attention management systems.

5.1 Attention Management vs. Ethics

Two directions drive the evolution of attention management systems. On one side, we have a major part of academic research, based on Weiser's vision of "*calm technology*" that for the most part is invisible to the user [116]. On the other side, there are commercial applications that compete for attention, sending over 100 notifications a day to individuals [69]. However, these two directions are not mutually exclusive. Commercial applications need to balance between grabbing a user's attention while avoiding to be excessively annoying. Similarly, "*polite*" attention management systems still need to ensure that information is delivered before it becomes stale. Consequently, designers should ensure that attention management systems comply with the following ethical restrictions: (i) attention fragmentation incurred by attention-seeking applications is held at an acceptable level and (ii) time-critical information is delivered in a timely manner.

Satisfying the latter restriction requires a deep knowledge of an individual's engagement, the relationship of the to-be-conveyed information and the current, as well as, future tasks a person is engaged in. In practice, the sender, or the application, might tag notifications with their urgency level. In version 8.0 of their mobile operating system Android, Google has introduced *notification channels*, a paradigm that allows applications to expose different streams of notifications separately so that users can fine-tune an individual channel's appearance. While this addresses notification differentiation within a single app, system-level priority settings for notifications suffer from the misuse of individual applications that may set all of their notifications to the highest priority. Practical means of ensuring that condition (i) is satisfied have been examined through a large part of this article – mobile sensing and machine learning harnessed to identify moments when a user is switching between tasks, is idle, or positively inclined towards the notification content, are the most commonly followed guidelines for designing "*calm*" attention management systems.

Another important ethical concern is related to the direction in which attention is steered. Interruptions might be welcomed by the user, the sentiment might be positive, and also the time to react might be low, yet, the overall effect on an individual's well-being or productivity may be negative. This is particularly important in case the application is used by vulnerable groups, such as children or those who suffer from a mobile phone or gaming addiction. However, we should be careful to disentangle technological determinism from the way technology is used – even when notifications are disabled on their phones, people still self-interrupt, checking for new content in fear of missing out [80].

Transparency. To ensure that the manipulation of human attention does not affect an individual's autonomy to select, accept, or ignore information, attention management systems should be designed transparently. Individuals need to have the option to *inspect*, *override*, or to *disable* functional principles. Interpretable attention management systems are one step in this direction, as they rely on human-readable rules to describe interruptibility preferences [67, 68]. A study by Mehrotra et al. indicates promising results towards transparent attention management systems, as more than 50% of generated rules were accepted by individuals [68]. Advances at operating system levels, such as *notifications channels*, are one step towards transparent notification management.

Privacy. Rich sensor data used to infer physical, physiological and mental states of individuals, information about the content of the message, relationship between the sender and receiver, have been widely used for efficient attention management. To ensure the privacy of both its users, as well as those interacting with users, attention management systems need to protect and safely process sensitive information. Such protection is especially important for systems that seek for a holistic view on an individual's interruptibility, gathering, exchanging, and storing information on multiple devices [77, 108]. For systems that construct interruptibility models from data harvested from more than one individual, *differential privacy* should be ensured, so that the probability of inferring individual behavior from data is minimized [34]. Finally, systems should incorporate best-practice privacy management approaches, such as the right for informational self-determination [47].

Research questions. Ethical concerns in ubiquitous attention management systems have not been widely addressed, yet. As guidelines for future efforts, we propose the following research questions:

(RQ1) *How should the chain of ethical responsibility be established and who or what are the key factors in it?*

(RQ2) *How should transparent attention management systems be designed? How should functional principles of attention management systems be presented to the user?*

(RQ3) *What should the requirements for privacy-aware attention management systems be? Which privacy regulations should be considered, as they may vary among legislators or organizations?*

(RQ4) *How to ensure that legal requirements are integrated within the technical development process while designing attention management systems?*

5.2 Evaluating Effects of Attention Management Systems on Human Well-being

Intelligent management of attention in ubiquitous systems may have a positive effect on human well-being [11, 51, 100, 121]. However, studies of such an effect are almost exclusively based on short time-scales, lasting approximately up to one month [100]. For example, Schneider et al. report short-term effects such as increased life balance, reduced stress levels, and reduced exhaustion of knowledge workers, after five-weeks of using mobile applications that manage availability [100]. Similar findings by Züger et al. suggest that, after two months of using their system, individuals feel less interrupted, more productive and experience more self-motivation [121].

Although studies show positive short-term effects, long-term effects of attention management systems on human well-being and cognition are rather unknown. The reasons for the lack of longitudinal studies might be the following: (i) additional burden on the users coming from the introduction of a new technology, (ii) reluctance of users to change their routine because of an intervention, and (iii) the challenge of acquiring large groups of participants for longitudinal studies. In fact, it has already been noted that the evaluation of attention management systems puts additional burden on individuals, as they have to be instructed to use new applications, fill out surveys regularly, or to be observed while using the evaluated solutions, which is the case if the *shadowing* evaluation method is used [41, 63]. Also, most attention management systems are designed as "*black boxes*". Without knowing functional principles, attention management restricts individuals' autonomy in selecting, accepting, and ignoring information. Individuals might resist the change these systems entail and try to restore settled routines [71]. Finally, large groups of participants are required for reliable long-term evaluations. To the best of our knowledge, the largest pool in this domain, with more than 680,000 participants (albeit for only three weeks and not measuring users' well-being), was obtained by Okoshi et al. in collaboration with Yahoo! Japan [78]. However, due to regulatory, privacy, and issues related to the additional burden imposed on participants, companies are in general reluctant to test prototype attention management systems' effect on well-being via production applications.

Towards Longitudinal Studies of the Effects of Attention Management Systems on Well-being. To facilitate long-term evaluation, we suggest that the following research directions are explored:

- **Unobtrusive well-being detection.** Stress and annoyance are the most commonly reported negative effects of interruptions in ubiquitous computing environments [1, 64]. ESM is the most common means of assessing users' psychological well-being, within which standardized tests, such as NASA-TLX, or simple Likert scales for emotion inference are often used. However, ESM puts additional burdens on individuals and should be used moderately, which limits the amount of data that can be collected. Recently, a new wave of pervasive devices – wearables – including smartwatches, fitness wristbands, and smart glasses brought physiological sensing to everyday use. Since physiological signals are correlated with certain aspects of well-being, sensor data streams from wearable devices could be used for unobtrusive well-being inference. For example, Gjoreski et al. use wrist-worn photoplethysmogram (PPG) sensors, which measure heart rate and blood pulse volume, to infer the stress level of a user [40]. The inference, however, is not straightforward, as PPG readings are affected by hand movement. Thus, accelerometers are used to detect whether a user is mobile or not before the stress level inference is performed. Besides stress, affects [101] and depression [79] can also be inferred unobtrusively thanks to mobile sensing.
- **Specific considerations in evaluation study design.** While attention management systems provide a technical solution to control interruptions, non-technical environmental factors may play a significant role in how a user's attention is influenced. For example, studies indicate that the design of workplaces, the awareness towards interruptions, or organizational measures have a direct impact on the number of interruptions [107, 121]. Therefore, efficient attention management systems need to consider such environmental factors. A holistic investigation of socio-technical aspects, including human, social, and organizational factors [12] might help define requirements for a long-term evaluation for a given environment. Without socio-technical requirements, individuals might perceive attention management systems as non-supportive, too complex to use [100], or even quit studies if they feel that a system does not reflect their interruptibility [121]. In one of our previous studies, users were reluctant to complete their study participation as the developed attention management tool did not support their well-established routine of using Skype for business communication [100]. Although, we focus on well-being in this section, the above concerns are relevant for all long-term evaluation studies of attention management systems.

We conclude the discussion with two questions meriting research:

(RQ6) How should a long-term evaluation of the well-being impact of attention management systems be performed provided minimal user burden and a large number of participants? The above discussion on unobtrusive well-being state detection could be a starting point for further exploration. Particularly, since physiological data can be incorporated in attention inference algorithms [43, 122].

(RQ7) How to incorporate contextual, social, and organizational factors when planning long-term evaluation studies? These aspects are not only important for evaluations but could also give essential features to the design of attention management systems.

5.3 Towards Theoretically-Grounded Attention Management Systems

Research on cognitive processes and ubiquitous computing systems practice have evolved separately, and one of the main goals of this survey is to identify untapped opportunities for building more efficient attention management systems that take into account the underlying psychological processes. In this subsection, we point out these potential research avenues, briefly discuss their theoretical underpinnings, and present work that, at least partly, addresses them.

Interrupting at task boundaries or idle times. Times, when a user is not actively engaged in a task, are the most suitable moments for interruption according to both theoretical frameworks [72], as well as experimental studies [1, 48]. From the memory for problem state theory point of view, at such moments the problem state, i.e., the intermediate mental representation necessary for performing a task, is not needed. Thus there is no interruption along procedural or declarative memory resources. Furthermore, task engagement flows through different stages: planning, execution, and evaluation [27]. In a study of instant messaging users, Czerwinski et al. found that the participants took more time to attend an interruption in case it arrived during the execution phase than in the other two phases [28]. Interestingly, they also find a disruptive effect of an interruption in the evaluation phase. The authors hypothesize that this may reflect the time required for users to visually re-orient themselves to where they left off and concluded the task completion. Attention management systems, such as Attelia II [77], successfully apply defer-to-breakpoint policies [50] within physical activities. Inferring one's mental engagement, on the other hand, is not straightforward. Certain physiological signals, such as pupil dilation, correlate with mental load, and indeed have been used for assessing task engagement [43]. However, a reliable means of detecting a user's task engagement using commodity devices is yet to be demonstrated [114]. Finally, Pielot et al. demonstrate that bored users are likely to engage with recommended content [85, 88].

(RQ8) How can we use commodity pervasive computing devices, such as wearable sensors, to infer a user's task engagement and consequently enhance attention management systems?

Problem state rehearsal. Bluma Zeigarnik's experiments from 1927 showed that people remember uncompleted or interrupted tasks better than completed tasks [120]. Therefore, there is a natural tendency to return to the primary task when interrupted. However, the complexity of the original task plays a significant role in the success of such a switchback. According to the Salvucci and Taatgen's theory, complex tasks are difficult to return to because they require the retrieval of the problem state from declarative memory. The more time the problem state spends in the declarative memory, the more likely it is that the retrieval will fail due to the decay of the stored information [98]. Borst et al. argue that the rehearsal of the problem state should come *before* a user switches to the secondary task [16]. Along the same lines, Cades et al. found that interruptions are more disruptive when they minimized the participant's ability to rehearse the primary task during the interruption, not necessarily when the secondary task is more difficult [19]. Trafton et al. have conducted experiments to demonstrate that visual alerts before a forced switch to another task help users rehearse the primary task state, and consequently facilitate the later switchback [110]. Currently, the rehearsal of the problem state represents an untapped challenge for practical attention management systems. Thus, we propose the following questions for future research:

(RQ9) How to condense and prepare the primary task for problem state rehearsal automatically? How to represent the content to an individual, so that the problem state is captured at a glance?

(RQ10) How long should the problem state rehearsal last, so that an interruption does not get stale?

Providing hints to revert the context. To facilitate a switchback, Salvucci and Trafton's theory suggest that cues related to the old problem state should be provided. Altmann and Trafton have experimentally confirmed that neighboring elements of the primary task are better retrieval cues than elements separated by greater psychological distance [5]. HCI researchers have been tackling this issue since the rise of multitasking PC applications. Czerwinski et al. have conducted a diary study of PC users to identify possible design improvements to facilitate task switchbacks. The authors propose time-centric visualizations and tools that reconfigure the layout of desktop windows (i.e., content and applications) that compromise a task [29]. Also, Iqbal and Horvitz examined the role of visual hints on task retrieval and found that the amount of visual obfuscation caused by an interruption impacts the time to return to the primary task – interruptions that cover only a small portion of a desktop are easier to return from [52]. Following this argumentation, attention management systems need to construct and provide psychologically close hints to support switchbacks to primary tasks. In [106], the authors

design a system for managing interruptions in intensive care units. They argue that color-coded notifications act as psychological cues that can be used to retrieve information about the patient it concerns. In this regard, the color of the notification helps to switchback to the patient *at a glance*. However, the automatic construction of psychological hints is still in its origin, and the following questions might help to steer future research in this area:

(RQ11) What are psychologically close hints and how can they be constructed automatically?

(RQ12) If multiple tasks can be resumed after an interruption, which hint should be displayed?

Relevance and relation to the primary task. Findings indicate that the interruption content plays an important role in the perceived disruptiveness of interruptions [105]. In fact, interruptions that contain relevant information for the primary task may negatively affect the time to complete the task but may help with the task performance quality. Interruptions with content irrelevant to the primary task also increase completion time but lead to higher task error rates [2]. A similar approach to relevance and relation of interruptions to the primary task is presented in [6]. The authors assume that an interruption is more relevant to individuals if it is matched with their social role (e.g., employee, family member). Consequently, attention management systems should incorporate the relevance of interruptions to the primary task. In this regard, we propose the following question:

(RQ13) How can an attention management system infer the relevance of an interruption to a task at hand, and at how should interruptions with different levels of similarity with the primary task be handled?

Multimodal alerts and notification representations. In ubiquitous systems, interruptions are often signaled through notifications that may use different alert types – sound, light, or vibration. Also, incoming information can be conveyed using different means – written on a screen, spoken through speakers, or by using a buzz alarm. The alert and message representation types may not align well with a user’s current task engagement. According to the cognitive theories presented in Section 3 different cognitive resources could be occupied at a time. If a user is engaged in a visual task only, perhaps an aural signal and notification whose contents are read to the user are still appropriate. An example would be a driving situation, where an incoming text message is less likely to distract the driver if read out aloud without requiring the driver to divert visual attention from the road. Consequently, a research question arises:

(RQ14) Can ubiquitous technologies infer which of the user’s cognitive resources are occupied and change the alert and information representation modality in a way to minimize disruption?

6 SUMMARY

In this article, we discussed the emerging field of attention management in ubiquitous environments. Particular attention has been paid to underlying theoretical frameworks and concepts of practical attention management systems as well as on actual implementations and technical designs. By leveraging unique characteristics of ubiquitous computing, e.g., personalized information, always-on connectivity or sensing capabilities, combined with cognitive and psychological concepts, attention management systems can help mitigate and mediate the ever-increasing number of interruptions. Fueled by theoretical and practical findings on the nature of human interruptibility, we discussed design implications for attention management systems, e.g., interrupting at task boundaries or providing cues for task retrieval. Actual implementations of such designs in the form of interruptibility models embedded in applications rely on mobile sensing, actuation, and machine learning. Although essential to efficiently support individuals in their attention management, building interruptibility models remains a challenge in the field of attention management. Identifying correlations between inferred contextual descriptors and an individual’s actual state of interruptibility represents an open research question. In this regard, we showed that attention management systems could benefit from interdisciplinary research as the facets of human interruptibility concern various research field, e.g., cognitive science, human-computer-interfaces,

or psychology. In fact, certain existing interruptibility models already incorporate findings from adjacent research fields, for example, interrupting at low cognitive states or between transitions of physical activities.

The capability to continuously sense, process, model, and to manage interruptibility facilitate adaptive attention management, more specifically, systems that can *sense* and *understand* the user. In the next few years, attention management will leverage existing and future modalities, e.g., smart clothes, sophisticated physiological measures, allowing for more efficient interruptibility models. Further investigations and efforts in this interesting field will help to shape and form efficient and advanced attention management systems.

REFERENCES

- [1] Piotr D. Adamczyk and Brian P. Bailey. 2004. If Not Now, when?: The Effects of Interruption at Different Moments Within Task Execution. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 271–278. <https://doi.org/10.1145/985692.985727>
- [2] Shamel Addas and Alain Pinsonneault. 2015. The many faces of information technology interruptions: a taxonomy and preliminary investigation of their performance effects. *Information Systems Journal* 25, 3 (2015), 231–273.
- [3] Erik M. Altmann and J. Gregory Trafton. 2002. Memory for goals: An activation-based model. *Cognitive science* 26, 1 (2002), 39–83.
- [4] Erik M. Altmann and J. Gregory Trafton. 2004. Task Interruption: Resumption Lag and the Role of Cues. *Proceedings of the Cognitive Science Society* 26, 26 (2004), 43–48.
- [5] Erik M. Altmann and J. Gregory Trafton. 2007. Timecourse of recovery from task interruption: Data and a model. *Psychonomic Bulletin & Review* 14, 6 (2007), 1079–1084.
- [6] Christoph Anderson, Clara Heissler, Sandra Ohly, and Klaus David. 2016. Assessment of Social Roles for Interruption Management: A New Concept in the Field of Interruptibility. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*. ACM, New York, NY, USA, 1530–1535. <https://doi.org/10.1145/2968219.2968544>
- [7] John R. Anderson. 1996. ACT: A Simple Theory of Complex Cognition. *American Psychologist* 51, 4 (1996), 355–365.
- [8] John R. Anderson, Mark V. Albert, and Jon M. Fincham. 2005. Tracing Problem Solving in Real Time: fMRI Analysis of the Subject-paced Tower of Hanoi. *Journal of cognitive neuroscience* 17, 8 (2005), 1261–1274.
- [9] Mark H. Ashcraft. 2006. *Cognition* (4 ed.). Pearson Prentice Hall, Upper Saddle River, NJ, USA.
- [10] GSM Association. 2018. *The Mobile Economy 2018*. Technical Report. GSM Association, London, United Kingdom. <https://www.gsma.com/mobileeconomy/>
- [11] Brian P. Bailey and Joseph A. Konstan. 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 (2006), 685–708. <https://doi.org/10.1016/j.chb.2005.12.009>
- [12] Gordon Baxter and Ian Sommerville. 2011. Socio-technical systems: From design methods to systems engineering. *Interacting with Computers* 23, 1 (2011), 4–17. <https://doi.org/10.1016/j.intcom.2010.07.003>
- [13] James Bo Begole, Nicholas E. Matsakis, and John C. Tang. 2004. Lilsys: Sensing Unavailability. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, New York, NY, USA, 511–514. <https://doi.org/10.1145/1031607.1031691>
- [14] Abraham Bernstein, Peter Vorburger, and Patrice Egger. 2008. A scenario-based approach for direct interruptibility prediction on wearable devices. *International Journal of Pervasive Computing and Communications* 3, 4 (2008), 426–438. <https://doi.org/10.1108/17427370710863149>
- [15] Christopher Bishop. 2006. *Pattern Recognition and Machine Learning* (1 ed.). Springer-Verlag New York, New York, NY, USA.
- [16] Jelmer P. Borst, Niels A. Taatgen, and Hedderik van Rijn. 2010. The Problem State: A Cognitive Bottleneck in Multitasking. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 36, 2 (2010), 363–382.
- [17] Jelmer P. Borst, Niels A. Taatgen, and Hedderik van Rijn. 2015. What Makes Interruptions Disruptive?: A Process-Model Account of the Effects of the Problem State Bottleneck on Task Interruption and Resumption. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2971–2980. <https://doi.org/10.1145/2702123.2702156>
- [18] Juliana J. Brixey, David J. Robinson, Craig W. Johnson, Todd R. Johnson, James P. Turley, and Jiajie Zhang. 2007. A Concept Analysis of the Phenomenon Interruption. *ANS. Advances in nursing science* 30, 1 (January 2007), 26–42.
- [19] David M. Cades, Deborah A. Boehm Davis, J. Gregory Trafton, and Christopher A. Monk. 2007. Does the Difficulty of an Interruption Affect our Ability to Resume? *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 51, 4 (October 2007), 234–238. <https://doi.org/10.1177/154193120705100419>
- [20] Linda L. Chao and Robert T. Knight. 1995. Human prefrontal lesions increase distractibility to irrelevant sensory inputs. *Neuroreport* 6, 12 (August 1995), 1605–1610.
- [21] Linda L. Chao and Robert T. Knight. 1998. Contribution of Human Prefrontal Cortex to Delay Performance. *Journal of Cognitive Neuroscience* 10, 2 (March 1998), 167–177.

- [22] Justin Cheng, Jaime Teevan, Shamsi T. Iqbal, and Michael S. Bernstein. 2015. Break It Down: A Comparison of Macro- and Microtasks. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 4061–4064. <https://doi.org/10.1145/2702123.2702146>
- [23] Marvin M. Chun, Julie D. Golomb, and Nicholas B. Turk-Browne. 2011. A Taxonomy of External and Internal Attention. *Annual Review of Psychology* 62, 1 (January 2011), 73–101. <https://doi.org/10.1146/annurev.psych.093008.100427>
- [24] Marvin M. Chun and Jeremy M. Wolfe. 2001. Visual Attention. In *Blackwell's Handbook of Sensation and Perception*, E. Bruce Goldstein (Ed.). Blackwell Publishing Ltd., Oxford, United Kingdom, 273–310.
- [25] Wesley C. Clapp and Adam Gazzaley. 2012. Distinct mechanisms for the impact of distraction and interruption on working memory in aging. *Neurobiology of Aging* 33, 1 (January 2012), 134–148. <https://doi.org/10.1016/j.neurobiolaging.2010.01.012>
- [26] Maurizio Corbetta and Gordon L. Shulman. 2002. Control of Goal-Directed and Stimulus-Driven Attention in the Brain. *Nature Reviews. Neuroscience* 3, 3 (March 2002), 201–215. <https://doi.org/10.1038/nrn755>
- [27] Edward B. Cutrell, Mary Czerwinski, and Eric Horvitz. 2000. Effects of Instant Messaging Interruptions on Computing Tasks. In *CHI '00 Extended Abstracts on Human Factors in Computing Systems (CHI EA '00)*. ACM, New York, NY, USA, 99–100. <https://doi.org/10.1145/633292.633351>
- [28] Mary Czerwinski, Edward Cutrell, and Eric Horvitz. 2000. Instant messaging: Effects of Relevance and Timing. In *People and computers XIV: Proceedings of HCI*, Vol. 2. ACM, New York, NY, USA, 71–76.
- [29] Mary Czerwinski, Eric Horvitz, and Susan Wilhite. 2004. A Diary Study of Task Switching and Interruptions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 175–182. <https://doi.org/10.1145/985692.985715>
- [30] Laura Dabbish, Gloria Mark, and Victor M. González. 2011. Why Do I Keep Interrupting Myself?: Environment, Habit and Self-interruption. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 3127–3130. <https://doi.org/10.1145/1978942.1979405>
- [31] Janez Demšar. 2006. Statistical Comparisons of Classifiers over Multiple Data Sets. *The Journal of Machine Learning Research* 7 (January 2006), 1–30.
- [32] Pedro Domingos. 2012. A Few Useful Things to Know About Machine Learning. *Commun. ACM* 55, 10 (October 2012), 78–87. <https://doi.org/10.1145/2347736.2347755>
- [33] J. Driver. 2001. A selective review of selective attention research from the past century. *British Journal of Psychology* 92, 1 (February 2001), 53–78.
- [34] Cynthia Dwork. 2006. Differential Privacy. In *Automata, Languages and Programming*, Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–12.
- [35] R. Fisher and R. Simmons. 2011. Smartphone Interruptibility Using Density-Weighted Uncertainty Sampling with Reinforcement Learning. In *Proceedings of 10th International Conference on Machine Learning and Applications and Workshops*, Vol. 1. IEEE Computer Society, Washington, DC, USA, 436–441. <https://doi.org/10.1109/ICMLA.2011.128>
- [36] James Fogarty, Andrew J. Ko, Htet Htet Aung, Elspeth Golden, Karen P. Tang, and Scott E. Hudson. 2005. Examining Task Engagement in Sensor-based Statistical Models of Human Interruptibility. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. ACM, New York, NY, USA, 331–340. <https://doi.org/10.1145/1054972.1055018>
- [37] Thomas Fritz, Andrew Begel, Sebastian C. Müller, Serap Yigit-Elliott, and Manuela Züger. 2014. Using Psycho-physiological Measures to Assess Task Difficulty in Software Development. In *Proceedings of the 36th International Conference on Software Engineering (ICSE 2014)*. ACM, New York, NY, USA, 402–413. <https://doi.org/10.1145/2568225.2568266>
- [38] R. Kelly Garrett and James N. Danziger. 2007. IM = Interruption Management? Instant Messaging and Disruption in the Workplace. *Journal of Computer-Mediated Communication* 13, 1 (October 2007), 23–42. <https://doi.org/10.1111/j.1083-6101.2007.00384.x>
- [39] Sonja Gievska and John Sibert. 2005. Using Task Context Variables for Selecting the Best Timing for Interrupting Users. In *Proceedings of the 2005 Joint Conference on Smart Objects and Ambient Intelligence: Innovative Context-aware Services: Usages and Technologies (sOc-EUSAI '05)*. ACM, New York, NY, USA, 171–176. <https://doi.org/10.1145/1107548.1107593>
- [40] Martin Gjoreski, Hristijan Gjoreski, Mitja Luštrek, and Matjaž Gams. 2016. Continuous Stress Detection Using a Wrist Device: In Laboratory and Real Life. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*. ACM, New York, NY, USA, 1185–1193. <https://doi.org/10.1145/2968219.2968306>
- [41] Victor M. González and Gloria Mark. 2004. "Constant, Constant, Multi-tasking Craziness": Managing Multiple Working Spheres. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 113–120. <https://doi.org/10.1145/985692.985707>
- [42] S. A. Grandhi and Q. Jones. 2009. Conceptualizing Interpersonal Interruption Management: A Theoretical Framework and Research Program. In *Proceedings of the 42nd Hawaii International Conference on System Sciences*. IEEE, Big Island, HI, USA, 1–10. <https://doi.org/10.1109/HICSS.2009.124>
- [43] Eija Haapalainen, SeungJun Kim, Jodi F. Forlizzi, and Anind K. Dey. 2010. Psycho-physiological Measures for Assessing Cognitive Load. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp '10)*. ACM, New York, NY, USA, 301–310. <https://doi.org/10.1145/1864349.1864395>

- [44] M. A. Hall and G. Holmes. 2003. Benchmarking Attribute Selection Techniques for Discrete Class Data Mining. *IEEE Transactions on Knowledge and Data Engineering* 15, 6 (Nov 2003), 1437–1447. <https://doi.org/10.1109/TKDE.2003.1245283>
- [45] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Advances in Psychology* 52 (1988), 139–183. [https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9)
- [46] Joyce Ho and Stephen S. Intille. 2005. Using Context-aware Computing to Reduce the Perceived Burden of Interruptions from Mobile Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. ACM, New York, NY, USA, 909–918. <https://doi.org/10.1145/1054972.1055100>
- [47] Gerrit Hornung and Christoph Schnabel. 2009. Data protection in Germany I: The population census decision and the right to informational self-determination. *Computer Law & Security Review* 25, 1 (2009), 84 – 88. <https://doi.org/10.1016/j.clsr.2008.11.002>
- [48] Shamsi T. Iqbal and Brian P. Bailey. 2005. Investigating the Effectiveness of Mental Workload As a Predictor of Opportune Moments for Interruption. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems (CHI EA '05)*. ACM, New York, NY, USA, 1489–1492. <https://doi.org/10.1145/1056808.1056948>
- [49] Shamsi T. Iqbal and Brian P. Bailey. 2006. Leveraging Characteristics of Task Structure to Predict the Cost of Interruption. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. ACM, New York, NY, USA, 741–750. <https://doi.org/10.1145/1124772.1124882>
- [50] Shamsi T. Iqbal and Brian P. Bailey. 2007. Understanding and Developing Models for Detecting and Differentiating Breakpoints During Interactive Tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 697–706. <https://doi.org/10.1145/1240624.1240732>
- [51] Shamsi T. Iqbal and Brian P. Bailey. 2008. Effects of Intelligent Notification Management on Users and Their Tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 93–102. <https://doi.org/10.1145/1357054.1357070>
- [52] Shamsi T. Iqbal and Eric Horvitz. 2007. Disruption and Recovery of Computing Tasks: Field Study, Analysis, and Directions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 677–686. <https://doi.org/10.1145/1240624.1240730>
- [53] N. Kern, S. Antifakos, B. Schiele, and A. Schwaninger. 2004. A Model for Human Interruptibility: Experimental Evaluation and Automatic Estimation from Wearable Sensors. In *Eighth International Symposium on Wearable Computers*, Vol. 1. IEEE, Arlington, VA, USA, 158–165. <https://doi.org/10.1109/ISWC.2004.3>
- [54] Nicky Kern and Bernt Schiele. 2006. Towards Personalized Mobile Interruptibility Estimation. In *Proceedings of the Second International Workshop on Location- and Context-Awareness*, Mike Hazas, John Krumm, and Thomas Strang (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 134–150. https://doi.org/10.1007/11752967_10
- [55] Kyohei Komuro, Yuichiro Fujimoto, and Kinya Fujita. 2017. Relationship Between Worker Interruptibility and Work Transitions Detected by Smartphone. In *Human-Computer Interaction. Interaction Contexts*, Masaaki Kurosu (Ed.). Springer International Publishing, Cham, 687–699.
- [56] Thomas Kubitz, Alexandra Voit, Dominik Weber, and Albrecht Schmidt. 2016. An IoT Infrastructure for Ubiquitous Notifications in Intelligent Living Environments. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*. ACM, New York, NY, USA, 1536–1541. <https://doi.org/10.1145/2968219.2968545>
- [57] Kostadin Kushlev and Elizabeth W. Dunn. 2015. Checking email less frequently reduces stress. *Computers in Human Behavior* 43, Supplement C. (February 2015), 220–228.
- [58] Kostadin Kushlev, Jason Proulx, and Elizabeth W. Dunn. 2016. "Silence Your Phones": Smartphone Notifications Increase Inattention and Hyperactivity Symptoms. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 1011–1020. <https://doi.org/10.1145/2858036.2858359>
- [59] Jennifer Lai, Sachiko Yoshihama, Thomas Bridgman, Mark Podlaseck, Paul Chou, and Danny Wong. 2003. MyTeam: Availability Awareness through the Use of Sensor Data. In *Proceedings of the International Conference on Human-Computer Interaction INTERACT*, Matthias Rauterberg, Marino Menozzi, and Janet Wesson (Eds.). IOS Press, Zurich, Switzerland, 503–510.
- [60] Sophie Leroy. 2009. Why is it so hard to do my work? The challenge of attention residue when switching between work tasks. *Organizational Behavior and Human Decision Processes* 109, 2 (July 2009), 168 – 181. <https://doi.org/10.1016/j.obhdp.2009.04.002>
- [61] Hong Lu, Denise Frauendorfer, Mashfiqui Rabbi, Marianne Schmid Mast, Gokul T. Chittaranjan, Andrew T. Campbell, Daniel Gatica-Perez, and Tanzeem Choudhury. 2012. StressSense: Detecting Stress in Unconstrained Acoustic Environments Using Smartphones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, New York, NY, USA, 351–360. <https://doi.org/10.1145/2370216.2370270>
- [62] Robert Malkin, Datong Chen, Jie Yang, and Alex Waibel. 2006. Multimodal Estimation of User Interruptibility for Smart Mobile Telephones. In *Proceedings of the 8th International Conference on Multimodal Interfaces (ICMI '06)*. ACM, New York, NY, USA, 118–125. <https://doi.org/10.1145/1180995.1181018>
- [63] Gloria Mark, Victor M. Gonzalez, and Justin Harris. 2005. No Task Left Behind?: Examining the Nature of Fragmented Work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. ACM, New York, NY, USA, 321–330.

- <https://doi.org/10.1145/1054972.1055017>
- [64] Gloria Mark, Daniela Gudith, and Ulrich Klocke. 2008. The Cost of Interrupted Work: More Speed and Stress. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 107–110. <https://doi.org/10.1145/1357054.1357072>
- [65] Daniel C. McFarlane. 1997. *Interruption of People in Human-Computer Interaction: A General Unifying Definition of Human Interruption and Taxonomy*. Technical Report No. NRL-FR-5510-97-9870. Office of Naval Research, Arlington, VA, USA.
- [66] Daniel C. McFarlane. 2002. Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. *Human-Computer Interaction* 17, 1 (2002), 63–139. https://doi.org/10.1207/S15327051HCI1701_2
- [67] Abhinav Mehrotra, Robert Hendley, and Mirco Musolesi. 2016. PrefMiner: Mining User’s Preferences for Intelligent Mobile Notification Management. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 1223–1234. <https://doi.org/10.1145/2971648.2971747>
- [68] Abhinav Mehrotra, Robert Hendley, and Mirco Musolesi. 2017. Interpretable Machine Learning for Mobile Notification Management: An Overview of PrefMiner. *GetMobile: Mobile Comp. and Comm.* 21, 2 (Aug. 2017), 35–38. <https://doi.org/10.1145/3131214.3131225>
- [69] Abhinav Mehrotra, Mirco Musolesi, Robert Hendley, and Veljko Pejovic. 2015. Designing Content-driven Intelligent Notification Mechanisms for Mobile Applications. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 813–824. <https://doi.org/10.1145/2750858.2807544>
- [70] Abhinav Mehrotra, Veljko Pejovic, Jo Vermeulen, Robert Hendley, and Mirco Musolesi. 2016. My Phone and Me: Understanding People’s Receptivity to Mobile Notifications. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 1021–1032. <https://doi.org/10.1145/2858036.2858566>
- [71] Anca M. Miron and Jack W. Brehm. 2006. Reactance Theory - 40 Years Later. *Zeitschrift für Sozialpsychologie* 37, 1 (Jan. 2006), 9–18. <https://doi.org/10.1024/0044-3514.37.1.9>
- [72] Yoshiro Miyata and Donald A. Norman. 1986. Psychological issues in support of multiple activities. In *User centered system design: New perspectives on human-computer interaction*, Donald A. Norman and S. W. Draper (Eds.). Lawrence Erlbaum, Hillsdale, NJ, USA, 265–284.
- [73] Christopher A. Monk, J. Gregory Trafton, and Deborah A. Boehm-Davis. 2008. The Effect of Interruption Duration and Demand on Resuming Suspended Goals. *Journal of Experimental Psychology: Applied* 14, 4 (December 2008), 299.
- [74] Darren Newton and Gretchen Engquist. 1976. The perceptual organization of ongoing behavior. *Journal of Experimental Social Psychology* 12, 5 (1976), 436–450. [https://doi.org/10.1016/0022-1031\(76\)90076-7](https://doi.org/10.1016/0022-1031(76)90076-7)
- [75] Tadashi Okoshi, J. Nakazawa, and H. Tokuda. 2014. Attelia: Sensing User’s Attention Status on Smart Phones. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication (UbiComp '14 Adjunct)*. ACM, New York, NY, USA, 139–142. <https://doi.org/10.1145/2638728.2638802>
- [76] Tadashi Okoshi, Julian Ramos, H. Nozaki, Jin Nakazawa, Anind K. Dey, and Hideyuki Tokuda. 2015. Attelia: Reducing User’s Cognitive Load due to Interruptive Notifications on Smart Phones. In *2015 IEEE International Conference on Pervasive Computing and Communications*. IEEE, St. Louis, MO, USA, 96–104. <https://doi.org/10.1109/PERCOM.2015.7146515>
- [77] Tadashi Okoshi, Julian Ramos, H. Nozaki, Jin Nakazawa, Anind K. Dey, and Hideyuki Tokuda. 2015. Reducing Users’ Perceived Mental Effort Due to Interruptive Notifications in Multi-device Mobile Environments. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 475–486. <https://doi.org/10.1145/2750858.2807517>
- [78] Tadashi Okoshi, Kota Tsubouchi, Masaya Taji, Takanori Ichikawa, and Hideyuki Tokuda. 2017. Attention and Engagement-Awareness in the Wild: A Large-Scale Study with Adaptive Notifications. In *2017 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, Big Island, HI, USA, 100–110. <https://doi.org/10.1109/PERCOM.2017.7917856>
- [79] V. Osmani. 2015. Smartphones in Mental Health: Detecting Depressive and Manic Episodes. *IEEE Pervasive Computing* 14, 3 (July 2015), 10–13. <https://doi.org/10.1109/MPRV.2015.54>
- [80] Antti Oulasvirta, Tye Rattenbury, Lingyi Ma, and Eeva Raita. 2012. Habits make smartphone use more pervasive. *Personal and Ubiquitous Computing* 16, 1 (01 Jan 2012), 105–114. <https://doi.org/10.1007/s00779-011-0412-2>
- [81] Fred Paas, Juhani E Tuovinen, Huib Tabbers, and Pascal WM Van Gerven. 2003. Cognitive load measurement as a means to advance cognitive load theory. *Educational psychologist* 38, 1 (February 2003), 63–71.
- [82] Veljko Pejovic and Mirco Musolesi. 2014. InterruptMe: Designing Intelligent Prompting Mechanisms for Pervasive Applications. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. ACM, New York, NY, USA, 897–908. <https://doi.org/10.1145/2632048.2632062>
- [83] Veljko Pejovic and Mirco Musolesi. 2015. Anticipatory Mobile Computing: A Survey of the State of the Art and Research Challenges. *Comput. Surveys* 47, 3 (April 2015), 1–29. <https://doi.org/10.1145/2693843>
- [84] Veljko Pejovic, Mirco Musolesi, and Abhinav Mehrotra. 2015. Investigating The Role of Task Engagement in Mobile Interruptibility. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15)*. ACM, New York, NY, USA, 1100–1105. <https://doi.org/10.1145/2786567.2794336>

- [85] Martin Pielot, Linas Baltrunas, and Nuria Oliver. 2015. Boredom-Triggered Proactive Recommendations. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15)*. ACM, New York, NY, USA, 1106–1110. <https://doi.org/10.1145/2786567.2794340>
- [86] Martin Pielot, Bruno Cardoso, Kleomenis Katevas, Joan Serrà, Aleksandar Matic, and Nuria Oliver. 2017. Beyond Interruptibility: Predicting Opportune Moments to Engage Mobile Phone Users. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 91 (Sept. 2017), 25 pages. <https://doi.org/10.1145/3130956>
- [87] Martin Pielot, Karen Church, and Rodrigo de Oliveira. 2014. An In-situ Study of Mobile Phone Notifications. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services (MobileHCI '14)*. ACM, New York, NY, USA, 233–242. <https://doi.org/10.1145/2628363.2628364>
- [88] Martin Pielot, Tilman Dingler, Jose San Pedro, and Nuria Oliver. 2015. When Attention is Not Scarce - Detecting Boredom from Mobile Phone Usage. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 825–836. <https://doi.org/10.1145/2750858.2804252>
- [89] Martin Pielot, Rodrigo de Oliveira, Haewoon Kwak, and Nuria Oliver. 2014. Didn't you see my message?: predicting attentiveness to mobile instant messages. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, New York, NY, USA, 3319–3328. <https://doi.org/10.1145/2556288.2556973>
- [90] Martin Pielot and Luz Rello. 2017. Productive, Anxious, Lonely: 24 Hours Without Push Notifications. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. ACM, New York, NY, USA, 11:1–11:11. <https://doi.org/10.1145/3098279.3098526>
- [91] Benjamin Poppinga, Wilko Heuten, and Susanne Boll. 2014. Sensor-Based Identification of Opportune Moments for Triggering Notifications. *IEEE Pervasive Computing* 13, 1 (January 2014), 22–29. <https://doi.org/10.1109/MPRV.2014.15>
- [92] Michael I. Posner. 1982. Cumulative Development of Attentional Theory. *American Psychologist* 37, 2 (February 1982), 168–179. <https://doi.org/10.1037/0003-066X.37.2.168>
- [93] Kiran K. Rachuri, Mirco Musolesi, Cecilia Mascolo, Peter J. Rentfrow, Chris Longworth, and Andrius Aucinas. 2010. EmotionSense: A Mobile Phones Based Adaptive Platform for Experimental Social Psychology Research. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp '10)*. ACM, New York, NY, USA, 281–290. <https://doi.org/10.1145/1864349.1864393>
- [94] Rahul Rajan, Ted Selker, and Ian Lane. 2016. Task Load Estimation and Mediation Using Psycho-physiological Measures. In *Proceedings of the 21st International Conference on Intelligent User Interfaces (IUI '16)*. ACM, New York, NY, USA, 48–59. <https://doi.org/10.1145/2856767.2856769>
- [95] Michael A. Regan, Charlene Hallett, and Craig P. Gordon. 2011. Driver distraction and driver inattention: Definition, relationship and taxonomy. *Accident Analysis & Prevention* 43, 5 (September 2011), 1771 – 1781. <https://doi.org/10.1016/j.aap.2011.04.008>
- [96] Claudia Roda and Julie Thomas. 2006. Attention aware systems: Theories, applications, and research agenda. *Computers in Human Behavior* 22, 4 (July 2006), 557–587. <https://doi.org/10.1016/j.chb.2005.12.005>
- [97] Dario D. Salvucci and Niels A. Taatgen. 2008. Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review* 115, 1 (January 2008), 101–130.
- [98] Dario D. Salvucci and Niels A. Taatgen. 2010. *The Multitasking Mind*. Oxford University Press, New York, NY, USA. 320 pages.
- [99] H. Sarker, M. Sharmin, A. A. Ali, Md M. Rahman, R. Bari, S. M. Hossain, and S. Kumar. 2014. Assessing the Availability of Users to Engage in Just-in-time Intervention in the Natural Environment. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. ACM, New York, NY, USA, 909–920. <https://doi.org/10.1145/2632048.2636082>
- [100] Katharina Schneider, Kathrin Reinke, Gisela Gerlach, Christoph Anderson, Sebastian Wojtek, Svenja Neitzel, Rahul Dwarakanath, Doreen Boehnstedt, and Ruth Stock. 2017. Aligning ICT-enabled Availability and Individual Availability Preferences: Design and Evaluation of Availability Management Applications. In *Proceedings of the Thirty Eighth International Conference on Information Systems*. InTech, Seoul, South Korea, 1–19.
- [101] Sachin Shah, J. Narasimha Teja, and Samit Bhattacharya. 2015. Towards affective touch interaction: predicting mobile user emotion from finger strokes. *Journal of Interaction Science* 3, 1 (11 Nov 2015), 6. <https://doi.org/10.1186/s40166-015-0013-z>
- [102] Saul Shiffman, Arthur A. Stone, and Michael R. Hufford. 2008. Ecological Momentary Assessment. *Annual Review of Clinical Psychology* 4, 1 (2008), 1–32. <https://doi.org/10.1146/annurev.clinpsy.3.022806.091415>
- [103] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. 2014. Large-scale Assessment of Mobile Notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3055–3064. <https://doi.org/10.1145/2556288.2557189>
- [104] Myeong-Ho Sohn, Adam Goode, V. Andrew Stenger, Kwan-Jin Jung, Cameron S. Carter, and John R. Anderson. 2005. An information-processing model of three cortical regions: evidence in episodic memory retrieval. *Neuroimage* 25, 1 (March 2005), 21–33. <https://doi.org/10.1016/j.neuroimage.2004.11.001>
- [105] Cheri Speier, Joseph S. Valacich, and Iris Vessey. 1999. The influence of task interruption on individual decision making: An information overload perspective. *Decision Sciences* 30, 2 (March 1999), 337–360. <https://doi.org/10.1111/j.1540-5915.1999.tb01613.x>

- [106] Preethi Srinivas, Anthony Faiola, and Gloria Mark. 2016. Designing Guidelines for Mobile Health Technology: Managing Notification Interruptions in the ICU. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 4502–4508. <https://doi.org/10.1145/2858036.2858553>
- [107] Edward R. Sykes. 2011. Interruptions in the workplace: A case study to reduce their effects. *International Journal of Information Management* 31, 4 (2011), 385 – 394. <https://doi.org/10.1016/j.ijinfomgt.2010.10.010>
- [108] Edward R. Sykes. 2014. A Cloud-based Interaction Management System Architecture for Mobile Devices. *Procedia Computer Science* 34 (August 2014), 625–632. <https://doi.org/10.1016/j.procs.2014.07.086>
- [109] Takahiro Tanaka and Kinya Fujita. 2011. Study of User Interruptibility Estimation Based on Focused Application Switching. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW '11)*. ACM, New York, NY, USA, 721–724. <https://doi.org/10.1145/1958824.1958954>
- [110] J. Gregory Trafton, Erik M. Altmann, Derek P. Brock, and Farilee E. Mintz. 2003. Preparing to resume an interrupted task: Effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies* 58, 5 (2003), 583–603.
- [111] Liam D. Turner, Stuart M. Allen, and Roger M. Whitaker. 2015. Interruptibility Prediction for Ubiquitous Systems: Conventions and New Directions from a Growing Field. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 801–812. <https://doi.org/10.1145/2750858.2807514>
- [112] Liam D. Turner, Stuart M. Allen, and Roger M. Whitaker. 2015. Push or Delay? Decomposing Smartphone Notification Response Behaviour. In *Human Behavior Understanding*, Albert Ali Salah, Ben J. A. Kröse, and Diane J. Cook (Eds.). Springer International Publishing, Cham, Switzerland, 69–83. https://doi.org/10.1007/978-3-319-24195-1_6
- [113] Liam D. Turner, Stuart M. Allen, and Roger M. Whitaker. 2017. Reachable but not receptive: Enhancing smartphone interruptibility prediction by modelling the extent of user engagement with notifications. *Pervasive and Mobile Computing* 40 (2017), 480–494. <https://doi.org/10.1016/j.pmcj.2017.01.011>
- [114] Gašper Urh and Veljko Pejovic. 2016. TaskyApp: Inferring Task Engagement via Smartphone Sensing. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*. ACM, New York, NY, USA, 1548–1553. <https://doi.org/10.1145/2968219.2968547>
- [115] Dominik Weber, Alexandra Voit, Philipp Kratzer, and Niels Henze. 2016. In-situ Investigation of Notifications in Multi-device Environments. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 1259–1264. <https://doi.org/10.1145/2971648.2971732>
- [116] Mark Weiser. 1991. The computer for the 21st century. *Scientific American* 265, 3 (September 1991), 94–104.
- [117] Svante Wold, Kim Esbensen, and Paul Geladi. 1987. Principal component analysis. *Chemometrics and Intelligent Laboratory Systems* 2, 1 (August 1987), 37 – 52. [https://doi.org/10.1016/0169-7439\(87\)80084-9](https://doi.org/10.1016/0169-7439(87)80084-9)
- [118] S. Yantis. 1998. Control of visual attention. In *Attention*, H. Pashler (Ed.). Psychology Press/Erlbaum (UK) Taylor & Francis, Hove, England, 223–256.
- [119] Fengpeng Yuan, Xianyi Gao, and Janne Lindqvist. 2017. How Busy Are You?: Predicting the Interruptibility Intensity of Mobile Users. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 5346–5360. <https://doi.org/10.1145/3025453.3025946>
- [120] Bluma Zeigarnik. 1927. The Retention of Completed and Uncompleted Activities. *Psychologische Forschung* 9 (1927), 1–85.
- [121] Manuela Züger, Christopher Corley, André N. Meyer, Boyang Li, Thomas Fritz, David Shepherd, Vinay Augustine, Patrick Francis, Nicholas Kraft, and Will Snipes. 2017. Reducing Interruptions at Work: A Large-Scale Field Study of FlowLight. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 61–72. <https://doi.org/10.1145/3025453.3025662>
- [122] Manuela Züger and Thomas Fritz. 2015. Interruptibility of Software Developers and Its Prediction Using Psycho-Physiological Sensors. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2981–2990. <https://doi.org/10.1145/2702123.2702593>
- [123] Sina Zulkernain, Praveen Madiraju, Sheikh Iqbal Ahamed, and Karl Stamm. 2010. A Mobile Intelligent Interruption Management System. *Journal of Universal Computer Science* 16, 15 (July 2010), 2060–2080. <https://doi.org/10.3217/jucs-016-15-2060>

Received November 2017; revised February 2018; accepted April 2018