Attention Computing: Overview of Mobile Sensing Applied to Measuring Attention

Aku Visuri
aku.visuri@oulu.fi
University of Oulu
Oulu, Finland

Niels van Berkel
n.vanberkel@ucl.ac.uk
University College London
London, United Kingdom

ABSTRACT

The measurement of participant attention is a frequent by-product of mobile sensing-based studies, which typically focus on user interruptibility or the effectiveness of notification deliveries. We note that, despite the popularity of interruptibility research within our discipline, research focused on attention is surprisingly scarce. This omission may be due to (a combination of) methodological, technological, or disciplinary constraints. In this paper, we argue how attention levels can be effectively measured with existing technologies and methodologies by adapting continuous measurements of attention fluctuations. Many clinically researched technologies, as well as sensing-based analysis methods, could be leveraged for this purpose. This paper invites co-researchers to assess the use of novel ways to measure attention in their future endeavours.

KEYWORDS

Attention computing, mobile sensing, attention, notifications

1 INTRODUCTION

Scientific inquiry into mobile sensing has predominantly assessed the effect of technology (*e.g.*, smartphones, smartwatches) on interruptiveness or related themes such as workflow or task effectiveness. Research on the widely used Experience Sampling Method (ESM), for example, attempts to interrupt participants at a suitable time to ensure that notifications are answered [11]. This type of research is often not labelled as a study on attention, but one could argue that it fits within the theme of attention computing. Similarly, when measuring work effectiveness or *workflow*, the outcomes of the study are often an assessment of the participant's level of attention towards their work task. There are earlier HCI works which investigate attention via methods like the dual-task paradigm [6] mapped with digital signals, but more recent works have moved towards using mobile sensing and more abstract - or at least, less specified - definitions of attention. We raise the question why these attention-related studies are typically not identified as such, whether by design, methodological challenges, or potentially simply due to the ambiguity of the term *attention*.

The research community has only recently embraced attention as a research theme. This is perhaps due to the lack of a clear definition or the absence of reliable measurement tools. After all, when we study themes empirically, it is nigh impossible to assess A's influence over B if we do not have concrete readings of the variables. Anderson et al. [1] begin their opening chapter of attention by stating that "*there is no common understanding of attention in the literature*". However, for the purpose of Human-Computer Interaction a clearly defined definition of attention can be beneficial; the selective concentration on a discrete stimulus (task) while ignoring other perceivable stimuli (*e.g.*, interruptions). Attention levels vary based on an individual's ability to withstand cognitive load – essentially the frequency and degree of non-task related stimuli presented to an individual.

Defining Attention

The terms *interruptibility* and *attention* are sufficiently distinct. Whereas interruptions consist of single, often brief, moments in time which only exist to influence the user, attention is a continuously tracked variable - only the *focus* of

ones attention varies. One can focus their attention either on a single entity [2] or share their attention between multiple entities [1]. A persons' level of attention can vary based on intrinsic or external attention focuses (*e.g.*, going shopping versus visiting the movies) [5].

When we measure interruptibility, what we are actually measuring is the degree with which an interruption diminishes our level of attention. The change in attention level, on the other hand, varies largely on our focus: how critical is the entity that has captured our attention. Moments when our attention is *not* focused on something are few and far between. Additionally, we can devote a very low level of attention towards our focus entities, and these focuses can shift very rapidly. Part of this behaviour is involuntary (*e.g.*, nothing to keep our mind busy), but it can also be something that can be influenced (*e.g.*, procrastination).

The definitions of attention as found in Psychology are partly abandoned when considering the perspective of Computer Science (CS). Our collaborative efforts as researchers are typically more geared towards interruptions – the external stimuli that exist to hinder our capabilities of maintaining attention – not to the level of attention itself. Although there is value in understanding when and how is the best way to interrupt users of technology, the current approach runs the risk of focusing too narrowly on the concept of interruption and thereby ignoring the broader concept of attention.

Introducing Mobile Sensing

But why do CS researchers not necessarily always leverage the Psychological definitions on attention? Perhaps the answer lies in the types of data we typically collect, and the difficulties encountered in mapping discrete signals (sensor values and subsequently interruptions) to something that is non-discrete by nature (attention). Another challenge lies in the challenge of measuring the user's attention. Asking the user to answer *e.g.* a self-report questionnaire using methods such as mobile Experience Sampling will by definition affect the participant's current level of attention, and signal mapping between sensor data and attention measurement is currently not feasible. Similar problems exists in affective computing, where human measurements still often rely on surveys and questionnaires which are either discrete ("moments in time") or summaries of longer periods of time, whereas contextual sensor data collection is continuous. Another ambiguity lies in attempting to dynamically understand how large the analysed time periods should be in order to accurately map sensed variables to any particular human behaviour. Uniform ideas of exactly how far into the past we should observe to understand behaviour that was exhibited (or reported) *now* simply do not exist.

The association between attention and interruptibility can work to our benefit, as well. The estimation of the interruptive nature of an entity can be reversed as a measurement of focus, or level of attention paid to the focus. Approaches like [13] are required to both deliver notifications with minimal interruptions, but also to collect maximum contextual data of where the users attention is focused. Oulasvirta et al. [10] showed that when moving, attention is mostly focused on paying attention to surroundings, only shifting to other tasks in short bursts. Thus, physical activity (including transportation methods) can be used as a strong indicator of attention availability. Similarly, Dingler et al. [4] showed that attention dwindles rapidly regardless of a task, but users can return to their attentive state within five minutes.

Challenges

The main critique towards attention computing this paper offers is the lack of continuous and dynamic attention sensing. Typically studies and models try to answer a binary question - should the user be interrupted *right now*, or not? This approach does not offer a more fine-grained assessment of the user's attentive state. Figure 1 visualises the way how attention level actually fluctuates over time as we shift our attention focus, and how we tend to identify breakpoints - moments we have concluded our attention towards an earlier focus and prepare for a new task requiring our attention. This paper identifies three different scenarios of attention shifts, but numerous others do exist:

- **Shifting focus - multitasking**: We are often required to rapidly shift our attention to multiple tasks which we work with in parallel - sometimes even simultaneously if it is physically feasible by e.g., using separate physical mediums (talking on the phone and driving, for example).
- **Wandering attention**: A scenario where we dont have a specific important task at hand, but our mind naturally wanders and brings different entities to our attention.
- **Ready state**: While maintaining a ready state and some form of attention towards an entity, we are expecting a more important task requiring our full attention to arrive in the near future.

Use of a binary classification not surprising, however, as the systems we build (and one could argue computers in general) are discrete systems. One approach used to circumvent this problem is bounded deferral, used in recent works like [4, 9]. The method involves postponing the delivery of an interruption to a more suitable time, which may or may not be defined at the moment of the initial assessment. This notion ties well to the theoretical model of attention, where the users attention towards a focus entity diminishes over
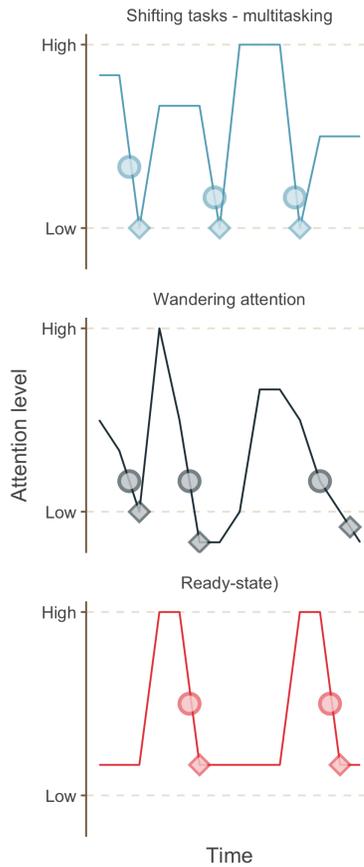
Figure 1: Abstract presentation of attention level fluctuation (y-axis) over time (x-axis). Square denotes a detectable *break-point* (shift of focus). Circle denotes a potentially more effective moment to attempt focus shift.

time - offering a moment in time at an undefined point of time in the future where the user's focus can be shifted to an interruption. Naturally, some interruptions can contain information which is simply deemed unnecessary by the user and never worth the users attention regardless of the user's current state [12].

Another, more open challenge, is defining attention and correctly labelling the user's attentive state. The models and methods from Psychology are rarely transformed into CS. This includes methods like measuring vigilance, cognitive load deterioration, reaction time, or the use of the dual-task paradigm. The methods we tend to use as CS and Ubicomp researchers tend to measure variables that share traits with attention - not attention directly.

Lastly, the challenges of capturing attention measurements *in-the-wild* are evident. As attention should be treated as a continuous shift throughout the day, simple laboratory studies do not capture these changes longitudinally. As with all modelling tasks, the primary challenge lies in labelling the data used to model attention throughout the day - especially in a fashion which could capture the continuous nature of attention, instead of discrete values at random points of the day.

2 MEASURING ATTENTION

The presented challenges of attention computing primarily focus on different methods for capturing attention. Thus far we have explored attention as a by-product of measurements like interruptibility (by e.g., calls and notifications), availability, workflow, and user's response to notification deferral schemes. This paper argues that much more precise methodologies exist, if we explore the psychological research spectrum in tandem with our own. Further cross-field studies exist where we could remove the by-product nature of contextual sampling and assess the attention measurements to notifications directly, for example.

In traditional neuropsychology, both vigilance and cognitive load deterioration are used as measurements of attention. One method employed consists of obtaining measurements on the subjects eye movement, and albeit the deployment of technologies like smartglasses appears halted, the research community is not hindered by such limitations and could use similar technologies to model and track attention conspicuously throughout the day. Numerous other biological signals can also be used to track the subject. Wearable sensors can use light to measure persons heart rate (HR), or blood pressure. HR and HR variability are often associated with stress, fatigue, and excitement, and these could be leveraged as significant factors in measuring attention. Similarly, small body worm devices can implement ECG (Electrocardiography) to measure heart activity or track oxygen saturation. The form factor of these devices can be a bracelet or even a ring. One recent example used electrodermal activity to measure student engagement in a classroom setting [3]. The Affectiva Q-Sensor [8] is scientifically validated and can track the subjects emotions, with the authors stating that "averages over longer time samples would be somewhat reliable". Similar more recent work by Hui et al. showcased an emotion sensing glove [7].

Simple device usage related metrics could also be refined to match the concept of attention in a more comprehensive level. Visuri et al. [12] showed that user availability towards ESM-type interruptions is low in the beginning of a usage session, but increases during the usage session. Interruption acceptance is significantly higher mere five minutes into a usage session. Meanwhile, Dingler & Pielot [4] showed that it takes five minutes for the users attentive state to return (for a particular task, messaging in this case).

This shows the fluctuating nature of attention, and even with these simple thresholds we could further investigate how ones attention changes during short time periods. The level of attention can also vary, but data such as application choice can provide us with insights on the user's task and attention level. Attention level could further reveal how (shape of the curve) and how quickly attention diminishes.

What prompts our device use? Unprompted device use can denote drops in attention, or at least shift in focus. Diurnal patterns - attention amplitude or rate of diminish over a period of a day - or even time window related patterns - are shifts in attention more likely to happen every 10 or 60 minutes? - could give much more detailed assessment of *continuous* tracking of attention than current ESM-based, more discrete, methodology. A significant portion of published research related to attention already includes vast details of user behaviour in matters that relate to attention - we could move from simply creating and testing new metrics to assess the effectiveness of work in other domains in the wild using the power of today's mobile devices.

Early approach used in attention management systems (in e.g., work environment) identified 'breakpoints'. As attention tends to diminish over time, we can further identify these breakpoints as moments when our attention level diminishes below a certain threshold, after which we feel like an interruption to our current focus is welcome. These breakpoints were often identified as proactively created by the subject, but a more detailed analysis of the subject's attention level fluctuations could even highlight scenarios where *introducing a breakpoint* (an interruption) could benefit the subject by resetting the attention level.

ACKNOWLEDGMENTS

3 CONCLUSIONS

This paper offers a preliminary discussion in how attention is an integral part of HCI research, while it is not always identified as such. We offer challenges, discussion points, and potential solutions and novel methods for increasing our ability to measure attention through related literature both within HCI as well as available technologies in other fields. We want to make measuring attention a goal by itself to ubicomp researchers - a goal that we are not afraid of but would rather embrace as a research goal. While exploring new methods will not be perfect from the get-go, studies could lead to suboptimal results, and mistakes could be made, if we do not try we will never succeed.

REFERENCES

[1] Christoph Anderson, Isabel Hübener, Ann-Kathrin Seipp, Sandra Ohly, Klaus David, and Veljko Pejovic. 2018. A survey of attention management systems in ubiquitous computing environments. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (2018), 58.

[2] MH Ashcraft. 2006. Cognition (4" ed.).

[3] Elena Di Lascio, Shkurta Gashi, and Silvia Santini. 2018. Unobtrusive Assessment of Students' Emotional Engagement During Lectures Using Electrodermal Activity Sensors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (2018), 103.

[4] Tilman Dingler and Martin Pielot. 2015. I'll be there for you: Quantifying Attentiveness towards Mobile Messaging. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services*. ACM, 1–5.

[5] Anja Exler, Marcel Braith, Andrea Schankin, and Michael Beigl. 2016. Preliminary investigations about interruptibility of smartphone users at specific place types. In *Proceedings of the 2016 ACM international joint conference on pervasive and ubiquitous computing: Adjunct*. ACM, 1590–1595.

[6] 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*. ACM, 301–310.

[7] Terence Hui and R Sherratt. 2018. Coverage of emotion recognition for common wearable biosensors. *Biosensors* 8, 2 (2018), 30.

[8] Arvid Kappas, Dennis Küster, Christina Basedow, and Pasquale Dente. 2013. A validation study of the Affectiva Q-Sensor in different social laboratory situations. In *53rd Annual Meeting of the Society for Psychophysiological Research, Florence, Italy*.

[9] Tadashi Okoshi, Kota Tsubouchi, and Hideyuki Tokuda. 2018. Real-world large-scale study on adaptive notification scheduling on smartphones. *Pervasive and Mobile Computing* 50 (2018), 1–24.

[10] Antti Oulasvirta, Sakari Tamminen, Virpi Roto, and Jaana Kuorelahti. 2005. Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile HCI. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 919–928.

[11] Niels van Berkel, Jorge Goncalves, Lauri Lovén, Denzil Ferreira, Simo Hosio, and Vassilis Kostakos. 2019. Effect of experience sampling schedules on response rate and recall accuracy of objective self-reports. *International Journal of Human-Computer Studies* 125 (2019), 118 – 128. https://doi.org/10.1016/j.ijhcs.2018.12.002

[12] Aku Visuri, Niels van Berkel, Tadashi Okoshi, Jorge Goncalves, and Vassilis Kostakos. 2019. Understanding Smartphone NotificationsâĂŹ User Interactions and Content Importance. *International Journal of Human-Computer Studies* (2019).

[13] Dominik Weber, Alireza Sahami Shirazi, and Niels Henze. 2015. Towards smart notifications using research in the large. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*. ACM, 1117–1122.