Investigating The Role of Task Engagement in Mobile Interruptibility

Veljko Pejovic
Faculty of Computer and Information Science,
University of Ljubljana
Slovenia
Veljko.Pejovic@fri.uni-lj.si

Mirco Musolesi
Department of Geography
University College London
UK
mirco.musolesi@ucl.ac.uk

Abhinav Mehrotra
School of Computer Science,
University of Birmingham
UK
axm514@cs.bham.ac.uk

MobileHCI '15 Adjunct, August 25–28, 2015, Copenhagen, Denmark.
ACM 978-1-4503-3653-6/15/08.
http://dx.doi.org/10.1145/2786567.2794336

Abstract

Context-awareness of mobile phones is a cornerstone of recent efforts in automatic determination of user interruptibility. Modalities such as a user's location, her physical activity, time of day, can be used in machine learning models to infer if a user is going to welcome an incoming notification or not. However, the success of context-aware interruptibility systems questions the existing theory of interruptibility, that is based on the internal state of the user, not her surroundings. In this work we examine the role of a user's internal context, defined by her engagement in the current task, on the sentiment towards an interrupting mobile notification. We collect and analyse real-world data on interruptibility of twenty subjects over two weeks, and show that the internal state indeed impacts user interruptibility.

Author Keywords

Interruptibility; Notifications; Multitasking; Context-aware computing.

ACM Classification Keywords

H.5.2. [Information Interfaces and Presentation (e.g. HCI)]: User Interfaces; H.1.2. [Models and Principles]: User/Machine Systems

Context-Aware Mobile Interruption Management

The mobile phone represents the most direct point of contact with almost any individual on the planet, and, since the smartphone revolution, also the most common means of getting the information from the Internet. Being every persons main communication link, the mobile often overwhelms its user by information. Consequently, efforts have been made to harness sensing and computing capabilities of modern mobiles to design smarter interruption systems that will not disturb users at times when they prefer to work uninterrupted.

The role that the context may have in determining user interruptibility has been recognised even before the smartphone era. In [3] Horwitz et al. build a Bayesian network that describes user-defined cost of interruption from contextual audio-visual features captured by a camera and a microphone in an office setting. Ho and Intille presented the first study of mobile interruptibility, in which wearable accelerometers were used to collect information about users' physical activity, while mobile notifications were sent at the same time [2]. Their study points out that the moments of changing physical activity represent preferred moments for interruption.

The smartphone has changed the way we reason about interruptibility. Its sensing and computing capabilities enable context-awareness previously explored only in so called Wizard of Oz studies [1]. Naturally, efforts have been made to infer user interruptibility from a richer physical context. Ter Hofte builds a model of the interruptibility of a smartphone user, that is based on a user's location, social circle and physical activity [9]. Yet, this work does not employ mobile sensing, but relies on self-reported information. Pielot et al. examine a communication- and usage-oriented context, by collecting and analysing a data set of text messages exchanged via smartphones together with the associated phone usage data [7]. Time since the screen was on, time since the last notification, and similar features were used in a classifier that infers if the users is going to attend the message within a short time frame. InterruptMe [6], on the other hand looks at the context in a more physical way. The system senses user's semantic location, physical activity and time of day, and connects this context with one's interruptibility through an evolving classifier.

Mobile Phone and the Theory of Multitasking

Not all interruptions are equally disruptive, and the importance of improving our understanding of disruptiveness was exemplified by the increasing amount of interruptions brought by the latest waves of communication technologies.

In the traditional theory of interruptibility three main factors that render an interruption disruptive have been identified: the interrupting task complexity, task duration, and the moment of interruption. A unifying theory of interruptibility is thoroughly explained in [8]. In a nutshell, our brain processes the current task in the procedural memory, while the declarative memory serves for the short- and long-term storage of facts. Once interrupted, the primary tasks is moved to the declarative memory, and the interrupting task takes over the procedural memory. Reactivation of the primary task after the interrupting task is finished may require recollection from the declarative memory. This recollection tends to be less successful if the interrupting task was complex and/or long. Time to rebuild the primary task's problem state also depends on the complexity of the primary tasks, and the familiarity of the subject with the primary task – the more practice one has with a certain task, the stronger the activation bound in the declarative memory will be. The moments of low cognitive workload, even the moments of break within a more complex task, tend to be more suitable

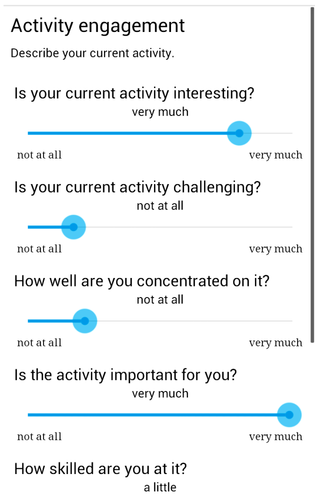

Figure 1: Screenshot of the ESM Android application. The survey was presented once a user clicked on a notification. Each notification was announced by the user's default notification ringtone and volume.

for interruptions, as shown in experimental studies [5].

However, from the mobile device point of view the estimation of the task complexity and the user's skills pertaining to the task, remain extremely challenging. Measuring the engagement of a user in a task, in order to identify low workload moments, has been demonstrated, with a limited success, via pupil dilation observation [5]. But the question we pose in this paper is: *Are task complexity and user engagement still relevant determinants of a user interruptibility in the mobile context?*

All the above studies have been performed in tightly controlled laboratory settings and without mobile devices. Furthermore, the theory of multitasking has been constructed on the data from unavoidable interruptions. In the mobile realm, interruptions are delivered via notifications. Here, in a so called bounded deferral manner [4], a user can decide to postpone the reaction until a more suitable moment. Given these differences, we decide to experimentally evaluate the relevance of task complexity and engagement on mobile interruptibility.

Experimental Setup

To experimentally investigate the relationship between task engagement and interruptibility, we develop and deploy an Android-based experience sampling method (ESM) application. The application occasionally delivers a notification on a user's phone; once the notification is clicked on, the user is presented with a brief survey, containing questions about the user's current engagement in a task. The questions include:

- Is your current activity **interesting**?
- Is your current activity **challenging**?
- How well are you **concentrated** on it?
- Is the activity **important** for you?
- How **skilled** are you at it?

Each question is answered on a four-point Likert scale, and possible answers include *not at all*, *a little*, *somewhat*, and *very much*.

In addition, a question *Is this a good moment to interrupt?*, answered on the same four-point Likert scale, is presented to a user in order to gauge her **sentiment** towards the interruption moment.

The data collection was carried out for two weeks among 20 adult subjects, who received a modest monetary reward for their participation, where "participation" was defined as having the application running on their phones, and no requirements were made on actually reacting to notifications. We recruited subjects from both sexes and from the age span 20 to 37 years old. Eight notifications were sent to each participant per day at random times between 8 am and 10 pm local time. In total 2334 notifications with surveys were sent, out of which 859 were fully answered, and our analysis concentrates on those answered messages.

Experimental Results

The goal of our study is not to build a fully contained predictive model of a mobile user's interruptibility and base it on the information about the user's current engagement in a task. Indeed, relying on the explicit user-provided information about task engagement, such a model would be highly impractical, while background sensing of this information may not be feasible with the current smartphone hardware. Instead, aim to investigate whether task engagement plays a role in mobile interruptibility, and if so, in which way different aspects of task engagement determine interruptibility.

To quantify the relationship between the five aspects of task engagement (here labelled as **Interesting**, **Challenging**, **Concentrated**, **Important** and **Skilled**) and the sentiment that the user expressed towards the moment of interruption, we fit a linear regression model with the sentiment as a dependent variable, and the five task engagement aspects as independent variables.

Table 1: Dep = Sentiment

Variable	Pearson Coefficient (Std. Err.)	t (Sig.)
Interesting	-.03 (.04)	-.82 (.410)
Challenging	-.12 (.03)	-3.87** (.000)
Concentrated	-.17 (.05)	-3.67** (.000)
Important	-.05 (.04)	-1.36 (.173)
Skilled	.17 (.04)	4.33** (.000)
(Regression Constant)	1.40 (.09)	15.21** (.000)
N	859	
R^2	0.10	
$F_{(5,854)}$	18.76 (p=.000)**	

The parameters of the fitted linear regression are shown in Table 1. The regression statistically significantly explains the sentiment towards interruptibility ($F(5, 854) = 18.76, p = .000$). However, the proportion of variation in the sentiment that can be explained by the engagement variables is not high ($R^2 = 0.10$), indicating that other factors that we have not considered in this analysis may influence the interruptibility.

Since all the survey questions are answered on the same Likert scale, we encode all the answers with the following values *not at all=0*, *a little=1*, *somewhat=2*, and *very much=3*. A positive correlation coefficient next to an independent variable indicates that a higher rating of this variable (i.e. from *a little* to *very much*) leads to a higher rating of the moment as a good moment to interrupt. In Table 1 only **Skilled** positively correlates with the sentiment. Thus, if a person is highly skilled in their current task, they are less likely to be irritated by an incoming notification. The other variables are negatively correlated with the sentiment. However, of those, only two – **Challenging** and **Concentrated** – are correlated at the significance level of 1% or less (denoted by **).

Discussion

Despite the change in the way we interruptions are communicated and negotiated with the user, the results from Table 1 are in accordance with the standard theoretical approach explaining interruptibility in a static context. For example, a user skilled in a task is less likely to be disturbed by an incoming mobile notification. A likely explanation is that such a user needs less time to reconstruct the primary task state after the interruption, due to stronger activation links in the declarative memory. On the other hand, the data shows that a user who is highly concentrated on a challenging task is more vulnerable to disruptions.

It is important to note that our metric of interruptibility is subjective – it is up to our users to declare their sentiment towards interruptibility. Since the study was done in an uncontrolled real-world setting, there was no way for us to measure the actual impact of an interruption on a primary task. It is also interesting that despite the possibility of deferring their actions for a more appropriate moment, users still feel annoyance when a notification arrives at times of high engagement. We suspect that the smartphone culture and social expectations around mobile usage require a user to keep information about a lingering notification in mind, even if she postpones her action and continues working on the original task.

Conclusion

Brain processes are in the foundations of models explaining human interruptibility. With the advent of personalised sensing devices, the opportunity opened to study interruptibility by sensing only the external context. Recent progress in that direction calls for re-examination of the need for inferring one's cognitive load in order to describe their interruptibility. Our study shows that although notifications allow users to defer interruptions for a later moment, the engagement with the current task still plays a significant role in determining users' interruptibility. For future work we leave automated recognition of task engagement via mobile sensors. In addition, we believe that understanding cognitive processes involved in multitasking can lead to further improvement of context-aware task management systems. For example, stimulating primary task recollection from the declarative memory via context-relevant links.

Acknowledgements

The analysis in this paper relies on the data collected as a part of the UK EPSRC-funded project “UBhave: ubiquitous and social computing for positive behaviour change" (EP/I032673/1).

REFERENCES

1. James Fogarty, Scott E. Hudson, Christopher G. Atkeson, Daniel Avrahami, Jodi Forlizzi, Sara Kiesler, Johnny C. Lee, and Jie Yang. 2005. Predicting Human Interruptibility with Sensors. *ACM Transactions on Computer-Human Interaction* 12, 1 (March 2005), 119–146.
2. Joyce Ho and Stephen S. Intille. 2005. Using Context-Aware Computing to Reduce the Perceived Burden of Interruptions from Mobile Devices. In *CHI'05*. Portland, OR, USA.
3. Eric Horvitz and Johnson Apacible. 2003. Learning and Reasoning about Interruption. In *ICMI'03*. Vancouver, Canada.
4. Eric Horvitz, Johnson Apacible, and Muru Subramani. 2005. Balancing awareness and interruption: Investigation of notification deferral policies. In *User Modeling 2005*. Springer, 433–437.
5. Shamsi T Iqbal and Brian P Bailey. 2006. Leveraging characteristics of task structure to predict the cost of interruption. In *CHI'06*. Montreal, Canada.
6. Veljko Pejovic and Mirco Musolesi. 2014. InterruptMe: designing intelligent prompting mechanisms for pervasive applications. In *UbiComp'14*. Seattle, WA, USA.
7. Martin Pielot, Rodrigo de Oliveira, Haewoon Kwak, and Nuria Oliver. 2014. Didn't You See My Message? Predicting Attentiveness to Mobile Instant Messages. In *CHI'14*. Toronto, ON, Canada.
8. Dario D Salvucci and Niels A Taatgen. 2010. *The multitasking mind*. Oxford University Press.
9. G. H. Ter Hofte. 2007. Xensible interruptions from your mobile phone. In *Mobile HCI'07*. Singapore.