Boredom-Triggered Proactive Recommendations

Martin Pielot
Telefonica Research
Placa Ernest Lluch i Marti, 5
Barcelona, 08019, Spain
martin.pielot@telefonica.com

Linas Baltrunas
Telefonica Research
Placa Ernest Lluch i Marti, 5
Barcelona, 08019, Spain
linas.baltrunas@telefonica.com

Nuria Oliver
Telefonica Research
Placa Ernest Lluch i Marti, 5
Barcelona, 08019, Spain
nuria.oliver@telefonica.com

 MobileHCI '15 Adjunct, August 25-28, 2015, Copenhagen, Denmark
ACM 978-1-4503-3653-6/15/08.
http://dx.doi.org/10.1145/2786567.2794340

Abstract
We propose the concept of boredom-triggered proactive recommendations for mobile phones. Given that more and more services attempt to attract mobile phone users' attention via push notifications, attention in this context can be considered as an increasingly scarce resource. However, when people are bored, by definition, they seek stimuli and hence often turn to their mobile phones. We cite evidence from our most recent work that boredom can be inferred from patterns of mobile phone usage and that during inferred phases of boredom, people are more likely to engage with suggested content. Thus, using boredom as content-independent trigger might help to make proactive recommendations a more pleasant experience and, in consequence, more successful.

Author Keywords
Boredom; Proactive Recommendations; Attention Economy; Notifications; Alerts; Availability; Interruptibility

ACM Classification Keywords
H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous

Attention and Proactive Recommendations
Human attention is valuable. The business model of two of the most dominant tech companies, namely Google and

Figure 1: Bored people often turn to their mobile phones. Source: "Show them this photo if someone said technology É" Adam Rifkin. May 21, 2014 via Flickr. CC BY 2.0.

Facebook, is primarily build around attention: they offer best-in-class internet services for free in exchange for the users' *eyeballs*, *i.e.* them paying attention to the contents of the services they offer. They monetize by selling the attracted attention to companies and individual who'd like to promote their content. The more time the users spent with the company's services, the more likely they are to focus fractions of their attention to the promoted content.

However, our attention is limited. The number of services that seek our attention can quickly deplete our attentional resources. Hence, economists, such as Davenport and Beck [4] propose to consider our attention as a *scarce resource* and coined the term *Attention Economy*. One of the next frontiers in this *battle for the user's attention*[1] are mobile devices. *Engagement is now defined by push-driven notifications rather than the traditional pull-driven experience*[2]. Recommendations will become proactive and notifications will be one essential path to deliver them.

Proactive recommendations are usually triggered by an event or change in a context. A recent work by Wang *et al.* showed [20] that the highest acceptance for push notifications could be achieved when both the content of the recommendation and the timing are well tuned. For example, *MailChimp*[3], an email sending service, has a functionality to automatically find the best performing hour of the day for delivering promotional emails. Wang *et. al* [19] used recent purchase events to trigger push notifications with advertisements for related products. Braunhofer *et al.* [1] used geo fencing to trigger music recommendations that fit a nearby Place of Interest.

Notifications are the Next Frontier

As with most scarce resources, we are facing the "tragedy of the commons" [8]: when individual companies behave rationally according to their self-interest by increasing their attempts to seek people's attention, they behave contrary to the best interests of the whole group by depleting the attentional resources of the user.

Furthermore, previous work has shown that notifications are associated with negative outcomes, such as decreased performance in the workplace [3], stress [10], and feelings of being overwhelmed and annoyed [14]. In particular notifications from non-communication applications are consid-

[1] http://www.wired.com/2014/12/new-media-2/
[2] http://techcrunch.com/2015/04/21/notifications-are-the-next-platform/
[3] http://eepurl.com/fLM42

ered of comparably low importance [18] and hence more likely to be dismissed or ignored [14].

Previous research has focused on identifying *opportune moments* for the delivery of notifications to mitigate their negative effects. Several works [9, 12, 13, 16, 17] show that *opportune moments* for different application contexts can be inferred from data available on the mobile phone. However, given the low importance that people assign to notifications from non-communication apps, seeking *opportune moments* may not be sufficient for proactive recommendations. While people may be willing to spare a few seconds to check a message from their loved ones, they may not be willing to devote attentional resources to interact and engage with a recommended item.

Attention is not Always Scarce

Yet, attention is not always scarce. One frequently occurring affective state [7] goes along with an abundance of attentional resources: *boredom*. Historically, boredom is defined as displeasure caused by "*lack of stimulation or inability to be stimulated thereto*" [6]. "*A bored person is not just someone who does not have anything to do; it's someone who is actively looking for stimulation but it is unable to do so*" [5]. Given these definitions, we assume that attention is abundant, rather than scarce, when a person is bored.

Mobile phones are a commonly used tool to kill time when people are bored [2]. In our most recent work on *When Attention is not Scarce: Predicting Boredom from Mobile Phone Usage* [15], we hypothesize that using the phone to kill time might result in distinct patterns of usage, which may be detected by a machine-learning model.

To investigate this hypothesis, we built an application called *Borapp* that collects objective and subjective data from the participants' mobile phones. Objective data is collected through the phone's sensors, *e.g.* location, time since last communication activity, or data-consumption rate. Subjective data was assessed via experience sampling. Approximately six times per day, a notification prompted users to respond to the statement: "*Right now, I am feeling bored.*" Responses were collected on a five-point Likert scale.

In a two-week deployment, we collected phone-usage data and 4398 valid self-reports from 54 participants. Our analysis showed that 35 features related to 7 different categories of phone usage could be used to infer boredom with high accuracy. The strongest predictors of boredom were related to these aspects: the recency of communication activity, intensity of phone usage (*e.g.* battery drain, number of apps launched, ...), the context (hour of the day and proximity sensor), and basic demographics.

From these features, we created a model to automatically infer the user's level of boredom. Our models achieve 74.6% to 82.9% AUCROC (Area Under the Receiver Operating Characteristic Curve) depending on how we formalize boredom – e.g. in terms of absolute responses or in terms of personalized variations in the responses – and whether we include data on the personal disposition towards boredom (boredom proneness [11]). We also experimented with the tradeoffs between precision and recall of boredom phases. When having 50% recall, the model achieved 62.6% precision, lowering recall to 30% resulted into 70.1% precision.

Boredom-Triggered Proactive Recommendations

This finding opens the door to using boredom as a signal for triggering proactive recommendations. Rather than requesting a person's attention at random and potentially generating the notification variant of *banner blindness*[4],

[4]Eye-tracking studies by the Normal-Nielssen Group shows that most people never watch the areas of web pages which

detected boredom could be used to limit proactive recommendations to moments when users would be more likely to appreciate them.

While the performance of the models were far from perfect, we proved in a follow-up study that their accuracy is sufficient to have measurable effects on the success rate of proactive recommendations. We updated *Borapp* to run the boredom-inference models online on the phone. The application collects the required data of the 35 features and infers boredom whenever one of the features changes. From time to time, the app triggered a proactive recommendation: a push notification suggested the user to read an article on Buzzfeed. Buzzfeed describes itself as the "*the most shareable breaking news, original reporting, entertainment, and video across the social web*"[5]. Thus, we assumed it might be reasonably interesting for most participants without being too popular. We tuned the recommendation engine to trigger notifications so that our data set would contain roughly 50% instances when the user is bored / not bored.

16 participants ran the application for two weeks. The results show that there were significant effects on the success of the proactive recommendations. Our participants were significantly ($p < .05$) more likely to open Buzzfeed and keep reading it for at least 30 seconds when the model inferred them to be bored. The ratio of opening Buzzfeed increased from 8% to 20% of the notifications, and the ratio of reading for 30 seconds or more increased from 4% to 15%. The effect size was large ($r > .500$) – hence, statistically speaking, the inference had a strong effect on the success-rate of the proactive recommendations.

Conclusions

In this position paper, we make the case for boredom-triggered proactive recommendations for mobile phones. This addresses the challenge that attention is scarce and that an increased volume of notification-delivered proactive recommendations may deplete people's attentional resources. We argue that this challenge needs to be addressed if proactive recommendations shall be widely successful.

As one solution, we suggest to use boredom, as a proxy for an emotional state where people actively seek stimuli, as a content-independent trigger for proactive recommendations. We cite evidence from our very recent work [15] that boredom can be inferred from mobile phone usage patterns and phone users, who are inferred to be bored, are significantly more open to suggested content.

Assuming that proactive recommendations delivered via mobile phone notifications will become more common in the future, using boredom as trigger will benefit service providers as well as the end users. End users will receive fewer recommendations that are triggered during times when they are busy. Service providers can use it to reduce the fraction of unsuccessful recommendations, which, for example, decreases the likelihood that users develop *banner blindness* towards proactive recommendations.

REFERENCES

1. Matthias Braunhofer, Marius Kaminskas, and Francesco Ricci. 2011. Recommending Music for Places of Interest in a Mobile Travel Guide. In *Proc. RecSys '11*. ACM.

2. Barry Brown, Moira McGregor, and Donald McMillan. 2014. 100 Days of iPhone Use: Understanding the Details of Mobile Device Use. In *Proc MobileHCI '14*. ACM.

contain the banners: http://www.nngroup.com/articles/banner-blindness-old-and-new-findings/

[5] http://www.buzzfeed.com/about

3. Mary Czerwinski, Eric Horvitz, and Susan Wilhite. 2004. A diary study of task switching and interruptions. In *Proc. CHI '04*. ACM.
4. Thomas H. Davenport and John C. Beck. 2002. *The Attention Economy: Understanding the New Currency of Business*. Harvard Business Press.
5. John D. Eastwood, Alexandra Frischen, Mark J. Fenske, and Daniel Smilek. 2012. The Unengaged Mind: Defining Boredom in Terms of Attention. *Perspectives on Psychological Science* 7, 5 (2012), 482–495.
6. Otto Fenichel. 1951. *On the psychology of boredom*. Columbia University Press.
7. Thomas Goetz, Anne C. Frenzel, Nathan C. Hall, Ulrike E. Nett, Reinhard Pekrun, and Anastasiya A. Lipnevich. 2014. Types of boredom: An experience sampling approach. *Motivation and Emotion* 38, 3 (2014), 401–419.
8. Garrett Hardin. 1968. The Tragedy of the Commons. *Science* 162, 3859 (1968), 1243–1248.
9. Eric Horvitz, Paul Koch, Raman Sarin, Johnson Apacible, and Muru Subramani. 2005. Bayesphone: Precomputation of Context-Sensitive Policies for Inquiry and Action in Mobile Devices. In *Proc UM '05*. ACM.
10. Shamsi T. Iqbal and Eric Horvitz. 2010. Notifications and awareness: a field study of alert usage and preferences. In *Proc. CSCW '10*. ACM.
11. Aleksandar Matic, Martin Pielot, and Nuria Oliver. 2015. Boredom-Computer Interaction: Boredom Proneness and the Use of Smartphone. In *Proc. UbiComp '15*. ACM.
12. Veljko Pejovic and Mirco Musolesi. 2014. InterruptMe: Designing Intelligent Prompting Mechanisms for Pervasive Applications. In *Proc. UbiComp '14*. ACM.
13. Martin Pielot. 2014. Large-Scale Evaluation of Call Availability Prediction. In *Proc. UbiComp '14*. ACM.
14. Martin Pielot, Karen Church, and Rodrigo de Oliveira. 2014. An In-Situ Study of Mobile Phone Notifications. In *Proc. MobileHCI '14*. ACM.
15. Martin Pielot, Tilman Dingler, Jose San Pedro, and Nuria Oliver. 2015. When Attention is not Scarce: Predicting Boredom from Mobile Phone Usage. In *Proc. UbiComp '15*. ACM.
16. Benjamin Poppinga, Wilko Heuten, and Susanne Boll. 2014. Sensor-Based Identification of Opportune Moments for Triggering Notifications. *Pervasive Computing* 13, 1 (2014), 22–29.
17. Stephanie Rosenthal, Anind K. Dey, and Manuela Veloso. 2011. Using decision-theoretic experience sampling to build personalized mobile phone interruption models. In *Proc. Pervasive '11*. Springer-Verlag.
18. Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. 2014. Large-scale Assessment of Mobile Notifications. In *Proc. CHI '14 (CHI '14)*. ACM.
19. Jian Wang, Badrul Sarwar, and Neel Sundaresan. 2011. Utilizing Related Products for Post-purchase Recommendation in e-Commerce. In *Proc. RecSys '11 (RecSys '11)*. ACM.
20. Jian Wang and Yi Zhang. 2013. Opportunity Model for e-Commerce Recommendation: Right Product; Right Time. In *Proc. SIGIR '13*. ACM.