Managing Smartphone Interruptions through Adaptive Modes and Modulation of Notifications

Hugo Lopez-Tovar
UCL Interaction Centre
University College London
h.tovar.11@ucl.ac.uk

Andreas Charalambous
Dept. of Computer Science
University College London
andreas.charalambous.10@ucl.ac.uk

John Dowell
UCL Interaction Centre
University College London
j.dowell@cs.ucl.ac.uk

ABSTRACT

Smartphones are capable of alerting their users to different kinds of digital interruption using different modalities and with varying modulation. Smart notification is the capability of a smartphone for selecting the user's preferred kind of alert in particular situations using the full vocabulary of notification modalities and modulations. It therefore goes well beyond attempts to predict if or when to silence a ringing phone call. We demonstrate smart notification for messages received from a document retrieval system while the user is attending a meeting. The notification manager learns about their notification preferences from users' judgements about videos of meetings. It takes account of the relevance of the interruption to the meeting, whether the user is busy and the sensed location of the smartphone. Through repeated training, the notification manager learns to reliably predict the preferred notification modes for users and this learning continues to improve with use.

Author Keywords

Adaptation; Smartphone notification; Interruption; Alert modality; Machine learning.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

Managing their digital interruptions is a pressing issue for smartphone users. All leading operating systems allow users to customize notification protocols and some third party apps considerably extend customization capabilities. But customized notifications are only a partial solution to the interruption issue. Apart from the overhead for users of learning to configure notification settings and continuing to maintain them [6], customized notifications do not discriminate between interruptions of varying importance and urgency in different contexts where users' receptivity varies [2]. Smart notification is the idea of an automatic personalisation of notifications, notifications adapted and optimised to the individual user. It would manage the user's digital interruptions either by selectively delaying notifications to a more suitable moment, or by selecting an appropriate modality or modulation for the notification, for example, by turning off alert sounds.

Smart notification has demonstrated that it is able to predict users' preferences for delaying phone call notifications in different contexts. The InterruptMe system [7] used explicit experience sampling to model a user's interruptibility in relation to context and then used this model to predict which notifications should be delayed. Approaches have also been demonstrated that mitigate the interruption caused by phone calls by silencing the notification but leaving the visual alert. The Ringlearn system [9] learned preferences for when to silence phone call notifications in different settings whilst the In-Context system [3] learnt when to silence phone call notifications taking account of different levels of user interruptibility.

The study of managing interruptions by smartphones follows the path of investigation in other contexts, for example performing office tasks with desktop computers [4, 8]. These studies have most often focused on managing phone calls, for example the Bayesphone system [5] demonstrated the use of the mobile device as a way to calculate interruption cost of incoming calls in the office context. However, phone calls are only one source of interruption for smart phone users who now must contend with notifications from a wide variety of digital services. These new kinds of interruption extend the vocabulary of notification with multiple modalities, formats and levels. For example, an interruption from a social network will often require a more nuanced notification than a phone call. Rather than simply a binary choice of alert now or later, or of let ring or silence [9], notifications require more carefully calibrated alert cues in visual, auditory and tactile modalities.

We demonstrate smart notification where the appropriate intervention is to adapt the notification modality and to modulate the alert cue. The notifications are messages from a simulated document retrieval agent system within a smart meeting room offering the user documents relevant to the meeting they are attending at particular junctures in the meeting. The system uses a Bayes Network to calculate the preferred notification method based on context where the data is obtained from participant analysis of videos of meetings. Evaluation showed the system to be capable of anticipating a majority of

users' notification preferences and will continue adapting to changes in the user's preferences.

STUDY AND DATA COLLECTION

To train the notification manager to predict a person's preferences for alerts while they are attending a meeting, we used recordings of meetings and asked participants to judge those meetings as though they themselves were present. We selected short videos of meetings [Fig. 1] from the AMI corpus [1] covering 22 different kinds of scenario. 70 participants were asked to watch the videos and for each they were asked to choose their preferred notification method given the context: imagining they were present in the meeting and either speaking or listening; supposing their mobile device was in a particular position (in their pocket; face down on the table; on a stand on the table; holding it) and expecting a given level of notification relevance (i.e. high/low). In this study, we have considered loud and discrete modulations for the sound modality. The rest of the modalities (i.e. text notification, flashing LED, vibration) don't allow different modulation levels.

The results showed that in only 45% of the scenarios would a binary classification of notification method (e.g. vibration and audible notification) be appropriate, while for the remainder, participants preferences were more or less equally distributed among the many available notification methods. Figure 2 shows examples for both situations and graphically confirms the argument for adapting mode and modulation of notifications to better calibrate the disturbance, rather than relying on a binary choice. Table 1 lists the summary of preferred notification methods by specific scenarios.

The full participants' responses have been aggregated to produce the default profile of preferences for our evaluation system, assuming it represents the notification preferences for the average user. This default profile of preferences consists of the probability distributions for each notification method given the context.

Figure 1. Example of meeting video presented to study participants. ©*Image from the AMI meeting corpus.*

How would you like to be notified if your phone is in your pocket?
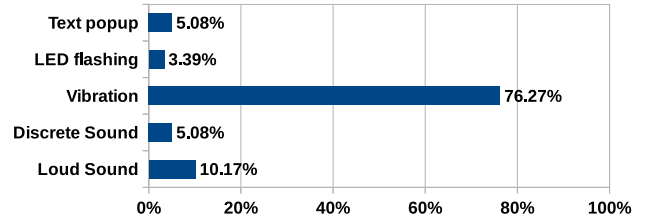
How would you like to be notified if your phone is in the table?
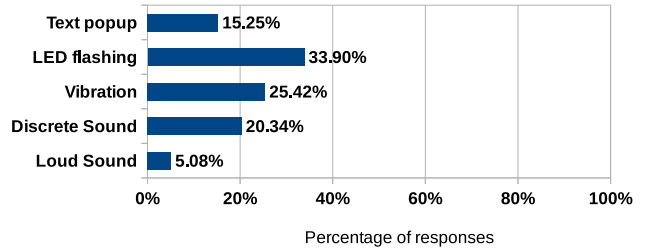
Percentage of responses

Figure 2. Above, a case where most of participants agree in a single notification method. Below, a case where participants' preferred method varies along the different notification methods.

ADAPTATION PROCEDURE

To infer the preferred notification mode given the context, a Bayesian Network is used as our statistical model, allowing us to represent the context as variables and their relationships as dependencies. The set of variables includes the relevance (low/high) of the interruption (i.e., whether the message from the document agent is currently relevant to the meeting), whether or not the person is speaking, and the mobile device position (pocket, stand, holding and table). Separate tests have shown that it is possible to reliably discriminate these different positions from a combination of accelerometer, gyroscope and ambient light sensor states.

The probability distributions of our model are initialised using the default profile generated from the collected data during our study and are updated with every feedback obtained from the user, i.e. whether they confirm that the proposed notification is acceptable. We have included a general probability distribution over the success of the notification methods considering only the device position (i.e. independent of notification relevance and speaking status) to overcome the problem of not having feedback on unseen scenarios, which is also updated with the participant responses.

Phone Loc.	High Relevance		Low Relevance	
In pocket	Vibr 71%	Vibr 76%	*Vibr 42%	*Vibr 47%
On the table	*LED 40%	*LED 34%	*LED 44%	LED 52%
On a stand	*LED 40%	*LED 49%	*LED 44%	*LED 45%
Holding it	Vibr 65%	Vibr 69%	*Vib 42%	*Vib/text 39%

Table 1. Summary of preferred notification methods from the initial collection of users' preferences, including the percentage of participants that selected the given notification method based on the notification relevance. Entries marked with * denote scenarios where other methods have also been selected by a significant number of participants.
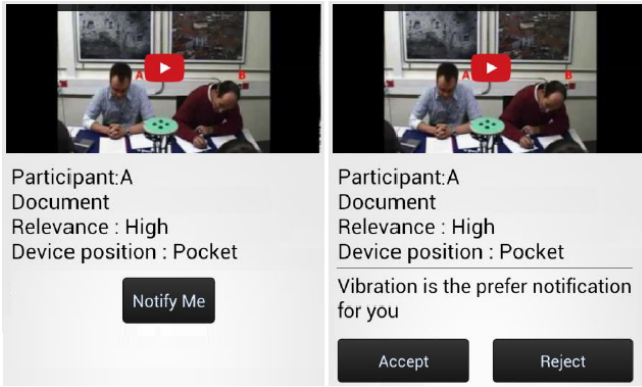
Figure 3. Left (3a), presented context scenario, allowing users to think about their preferred notification method. Right (3b), inferred notification method which can be accepted or rejected. ©*Images from the AMI meeting corpus.*

In the case of an accepted notification, the probability distributions are updated to give advantage to this particular notification method given the variable values. When the participant rejects the notification mode, their response doesn't include information about their actual preferred mode, so in this case the probability distributions are adjusted to reduce the likelihood of this particular notification method, favouring the others. To allow the agent to adapt to the participant's current preferences, we have implemented an ageing process to give more weight to recent user feedback. We calculate the weight w by a decreasing function of the form $w = k(x^2)$ where x is the number of days a decision-response pair was given as feedback by the user, and k is a constant inversely proportional to the maximum number of days being considered.

EVALUATION

Setup
To evaluate our prototype smart notification manager, ten participants were asked to confirm its predictions about their preferred notifications, again using a set of meeting videos. At this stage of course the Bayesian Network model was using the default profile generated from the study data and so its predictions were for a nominal average user. Its ability to adapt to the personal preferences of individual participants would be a key criterion for the system's success.

For the system's training phase, participants were presented with a set of short meeting videos representing different scenarios. For each video, the participants would first be presented with the context information, allowing them to picture themselves in the given situation [Fig. 3a]. They were told which person they should consider themselves to be in the video (e.g. person A), how relevant to the meeting is the document being offered, and the current position of their smartphone. The notification manager then proposes to them the method of notification it predicted they would prefer in this precise situation (e.g., "Vibration is the preferred notification for you"). The participant would then confirm whether this prediction was correct [Fig. 3b]. The feedback obtained from each scenario was used to update the probability distributions
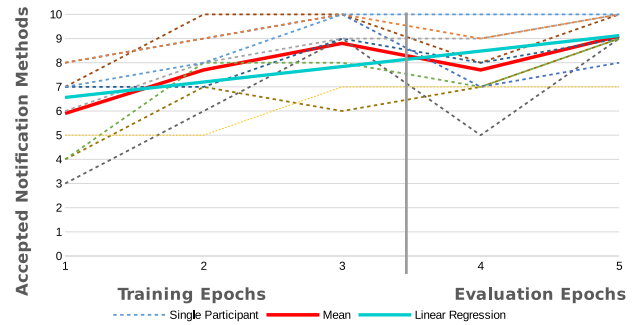

Figure 4. Number of accepted inferred notification methods by epoch. Each dotted line represents a single evaluation participant, solid lines represent the average acceptance and the linear regression.

for the next videos. To evaluate the progress of adaptation, we grouped the presented videos in sets of 10, each group constituting an 'epoch'. For the training phase, we presented the same 10 videos during three epochs to evaluate the adaptation on previously seen videos. We then presented two more epochs of 10 unseen videos each to evaluate adaptation on unseen scenarios. This made a total of 30 videos (i.e. 10 for training, 20 for evaluation).

Results
Figure 4 shows a plot with the evaluation results by epoch, indicating that:

- By the end of the training phase, every participant accepted from 2 to 6 more proposed notifications methods.

- By the end of the evaluation phase, most participants accepted from 2 to 6 more proposed notifications methods than before adaptation, except one which began and ended in 80% acceptance (reaching 100% at the end of training phase).

By analysing the mean of the participants' responses, we can see that:

- By using the default profile (i.e. no adaptation), the initial notification method was accepted in 5.92 out of 10 scenarios with a dispersion in the range of 3 to 8.

- By the end of the training phase, the number of accepted notification methods had increased to an average of 8.8 in 10 scenarios (a 60% improvement).

- By the end of the evaluation phase, the number of correctly predicted notification methods had increased to 9.1 (a 70% improvement compared to no adaptation).

Using linear regression analysis, we calculate a Correlation Coefficient r of 0.8056, confirming strong positive correlation between the epoch and the amount of accepted notification methods. In other words, as the number of epochs increase, more notification methods are accepted. Analysis of variance over all the epochs shows that the number of accepted notifications increases significantly ($F = 7.587$, $P < 0.0001$).

These results show that the default profile is an appropriate starting point and that the system continues to adapt to the

user's preferences in an agile way. They also demonstrate that it is possible to adapt by exploiting feedback limited to acceptance or rejection, without requiring the user to provide their actual preferred notification method.

CONCLUSIONS

Smartphones are capable of providing notifications in many different ways for different kinds of situation. That capability is under-exploited and typically only a manual switching between ringing and vibrating notifications for phone calls is used. The idea of smart notification is of a smartphone able to learn a user's preferences for notifications of different kinds in different situations. Notification would adapt to the type of interruption, the user's receptivity to interruptions and the particular context including where the device itself is currently positioned.

Research on smart notification has to date focused narrowly on managing phone call notifications. We have investigated document delivery notifications by smartphones in the setting of a smart meeting room. We have confirmed the need for adapting notifications to multiple modalities and modulations through a study where participants were asked to watch short videos of meetings and to identify their preferred notification method given the context as though they themselves were present in the meeting. Results show that in slightly more than half of the scenarios, users' preferred notification methods cannot be categorised only in two modalities.

For our adaptation model we chose a Bayesian Network that allows us to define the context as variables and their relationships as dependencies. The Bayesian Network is initialised with the probability distributions of the consensus data collected from the study. This proved to be a suitable starting point for our model which reached an average of 59% of accepted notifications before adaptation took place, in a range of 30% to 80%, which shows the dispersion of users' preferences.

For the evaluation of our model the results show that the system continuously adapts to the user's preferences in an agile way for both already experienced and unseen scenarios. Results also show that acceptance/rejection feedback is enough for our system to adapt, easing the work to the users.

The presented work contributes to research on smart notification by:

- expanding the scope of work to exploit the tasks and events that smartphones now make available, beyond simply incoming calls.

- demonstrating that broadening the vocabulary of notification modalities and modulations improves adaptation to manage users' interruptions by their smartphones.

Although conducting the validation in the laboratory using meeting videos rather than in a trial involving real meetings does not fully validate the efficacy of the system, it provides nevertheless quite compelling evidence. Whilst a field trial would offer much greater ecological validity, the lack of controls would mean that the effects of the adaptation process we have developed would be observed together with effects from other contextual factors, hence the importance of clear results in the laboratory. Having established these results in the laboratory, a field trial of the system should now be conducted.

In our model, our interpretation of users' interruptibility is limited to whether or not they are talking. Broadening this concept, as well as the representation of the context, could help to improve the adaptation to more specific situations. It also remains to investigate to what degree users are tolerant of adaptation inaccuracy.

REFERENCES

1. Carletta, J., Ashby, S., Bourban, S., Flynn, M., Guillemot, M., Hain, T., Kadlec, J., Karaiskos, V., Kraaij, W., Kronenthal, M., et al. The ami meeting corpus: A pre-announcement. In *Machine learning for multimodal interaction*. Springer, 2006, 28–39.

2. Fischer, J. E., Yee, N., Bellotti, V., Good, N., Benford, S., and Greenhalgh, C. Effects of content and time of delivery on receptivity to mobile interruptions. In *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*, ACM (2010), 103–112.

3. Fisher, R., and Simmons, R. Smartphone interruptibility using density-weighted uncertainty sampling with reinforcement learning. In *Machine Learning and Applications and Workshops (ICMLA), 2011 10th International Conference on*, vol. 1, IEEE (2011), 436–441.

4. Horvitz, E., and Apacible, J. Learning and reasoning about interruption. In *Proceedings of the 5th international conference on Multimodal interfaces*, ACM (2003), 20–27.

5. Horvitz, E., Koch, P., Sarin, R., Apacible, J., and Subramani, M. Bayesphone: Precomputation of context-sensitive policies for inquiry and action in mobile devices. In *User Modeling 2005*. Springer, 2005, 251–260.

6. Khalil, A., and Connelly, K. Context-aware configuration: A study on improving cell phone awareness. In *Modeling and Using Context*. Springer, 2005, 197–209.

7. Pejovic, V., and Musolesi, M. Interruptme: designing intelligent prompting mechanisms for pervasive applications. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ACM (2014), 897–908.

8. Shrot, T., Rosenfeld, A., Golbeck, J., and Kraus, S. Crisp: an interruption management algorithm based on collaborative filtering. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, ACM (2014), 3035–3044.

9. Smith, J., and Dulay, N. Ringlearn: Long-term mitigation of disruptive smartphone interruptions. In *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2014 IEEE International Conference on*, IEEE (2014), 27–35.