Sensor-based Identification of Opportune Moments to Trigger Unobtrusive Notifications

Benjamin Poppinga, OFFIS – Institute for Information Technology, Oldenburg, Germany, benjamin.poppinga@offis.de

Wilko Heuten, OFFIS – Institute for Information Technology, Oldenburg, Germany, wilko.heuten@offis.de

Susanne Boll, University of Oldenburg, Germany, boll@informatik.uni-oldenburg.de

Attention: This is a pre-print version and therefore slightly different from the final published article.

Abstract

Notifications are important means to attract the user's attention on a mobile device, e.g., if a text message is received. Today's smart phones issue the notification immediately after an event occurred and then repeat unanswered notifications in fixed time intervals. The disadvantage of this issue and repeat strategy is that notifications can appear in inconvenient situations and are therefore perceived as annoying and interrupting. In this paper, we study what the mobile context as inferred through a phone's sensors for both, answered and ignored, notifications looks like. To do so we conducted a large-scale, longitudinal study via the Google Play store and observed 6581 notifications from 79 different users over up to 76 days. A derived model is able to predict opportune moments to issue notifications with a reasonable accuracy of about 77 %. We argue that our findings can lead to intelligent strategies to issue unobtrusive notifications on today's smart phones at no extra cost.

Keywords

Interruption, Attention, Notification, Mobile, User Modeling, Evaluation

Introduction

Today's smart phones often use notifications to attract the user's attention on the device to indicate that something more or less important has happened. Typical reasons for a notification are, e.g., an incoming text message or a set reminder. While some notifications have informatory reasons, e.g., an available application update, and can be handled with delay, most notifications require immediate user attention and action. For this reason, most notifications are presented in an obtrusive way, e.g., by visual screen appearance, short vibration, and by flashing an LED. If not attended, some notification reminders are repeated.

It is undoubtedly valuable to have means like notifications to quickly reach the user. However, the existing *issue and repeat* strategy can be quite obtrusive and annoying as notifications and their repetitions might occur in inconvenient situations like at night time or when driving a car. These unwanted notifications often lead to stress, increased frustration, time pressure and effort [1]. Creating context-awareness by modeling the user's interruptibility has been proposed to be a solution to this problem [2]. The idea is that a context-aware phone somehow identifies if it is appropriate to trigger a notification and thereby reduce obtrusiveness. While it has been shown that such opportune and non-disruptive moments can be reliably identified for workers on stationary desktop computers in [2, 3], research is still ongoing for mobile and dynamic outdoor scenarios.

In this paper, we investigate in which contexts mobile notifications are typically attended or not. To study this behavior we designed MoodDiary, a mobile diary application for mood tracking, and distributed it via the Google Play store. A resulting model, which we derived from our collected 6581 total notifications from 79 users, is 77.85 % accurate in predicting opportune moments to issue notifications. In contrast to earlier work, our observations are made in real life, *in the wild* and over a long period of time, which results in high external validity and broad applicability of our findings. We argue that our insights can help designers and developers of mobile software to trigger notifications in less disruptive moments. Users will probably benefit from stress reduction and a productivity boost.

Related Work

Interruptions distract humans from their ongoing activity by introducing new tasks. It is differentiated between internal interruptions, i.e., those that appear in our own thought process, and external interruptions, those that have their cause in the environment [4].

With technological advance there are more and more machines touting for our attention, causing external interruptions. Because machines are insensitive to whether a human is able to attend a notification or not [5], active interruption management is needed. In theory four design solutions to cope with interruptions were identified: immediate, negotiated, mediated, and scheduled interruption [6]. Mediated interruptions use indirect information, e.g., a human's digital calendar or environmental sensors, and are recognized among the less interruptive techniques [7].

While there are many factors that influence a person's interruptibility, e.g., social engagement, literature found the user activities to be very relevant. Interruptions between coarse breakpoints, i.e., major changes in the work flow, produce less annoyance and users assessed them as being more respectful to their ongoing activity compared to other, e.g., random moments [8, 9]. To detect and differentiate between these breakpoints in interactive tasks, models can be created. Such models for stationary computers in office contexts reach an accuracy of up to 78 % [3].

However, in comparison to stationary computers, mobile devices are used in much more dynamic ways and in by far more complex contexts. Fischer et al. argue that the endings of mobile interactions, like making phone calls and receiving SMS, denote a coarse breakpoint and are opportune moments to attend notifications. They found that a user deals with a notification more quickly if it is triggered at coarse breakpoints. Further, they found that the type of task has a significant effect on acceptance time and completion rate. Despite the triggering in coarse breakpoints, some users still noted that notifications and subsequent tasks were annoying in certain situations [10].

Instead of using user-entered and phone-specific information, like a received SMS, also sensors could be used to identify opportune moments. Ho and Intille used two accelerometers and activity recognition to trigger interruptions at moments when user changes activity, e.g., from sitting to walking [11]. They showed that users are significantly more receptive to interruptions compared to a random condition. Kern and Schiele used wearable accelerometers, audio, and location sensors to decide whether a user should be notified and which modality to use. They argue that no advanced model is needed, as the combination of tendencies is already sufficient [12].

In contrast to earlier research, we ground our research on sensors and technologies that are available in today's commercial mobile phones. Earlier work [3, 10, 11] studied about 20 users for a maximum of 14 days, whereby they investigated a maximum of 2000 notifications and achieved quite artificial response rates of up to 90 %. In contrast, we investigate the users' behavior in a large-scale, longitudinal study. We studied 79 users over up to 76 days and investigated 6581 notifications. Thereby we reach a practical response rate of about 23 %. We argue that our findings have a high external validity and their application can lead to a significant reduction in notification obtrusiveness.

Concept

In this paper, we study the contexts in which notifications are considered and in which they are rejected or ignored by the user. To do so we design an application, which logs the context a user is in when a notification is triggered and if this notification is answered or rejected. We will use the gathered data to develop a predicting model on what are opportune moments to issue a notification in mobile contexts.

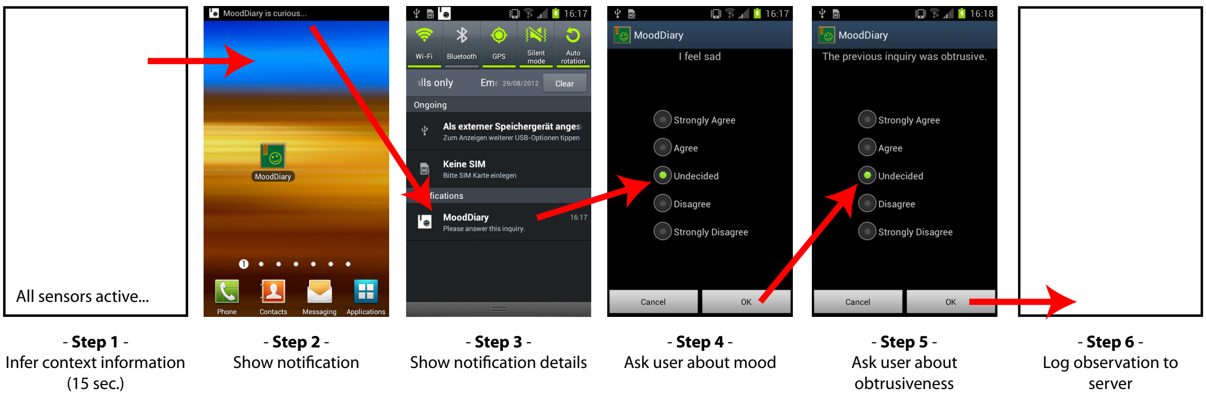

Figure 1: The process how a notification is issued and answered consists of six steps. In each of the steps that require user interaction a notification can be rejected. In both cases, answering and rejection, the context information are transferred to our servers (Step 6).

Context Observation Approach

A critical decision to make is which information should be considered in the to-be-created predictive model. Previous work used either user-entered information, like calendar entries, or sensor data. User-entered information has heterogeneous levels of availability and quality, as it might be sloppily maintained and is therefore incorrect or outdated. Further, this information is personal and users might not want to share it with an unknown application or third party. In contrast, sensor data is homogeneous as most smart phones come with similar sensing capabilities and the measured values are often normalized and well defined.

Consequently, we follow the sensor-based context observation approach as we think this is more reliable and promising. Our approach incorporates sensors which are available in nowadays smart phones: GPS, accelerometer, gyroscope, compass, microphone, and a proximity sensor. These sensors can obtain various measures, which can be used to observe various aspects of the users' notification answering behavior. In fact, the recorded context features are: timestamp, location provider, position accuracy, speed, GPS heading, compass heading, roll, pitch, proximity, and light level. We neither collected location information nor assessed the microphone for ethical reasons. Although we use state of the art sensors, we want to emphasize that these can only observe a tiny fragment of what makes a holistic mobile context.

Apparatus: MoodDiary

To collect context information and trigger notifications on the users' mobile phones, we needed an apparatus. For our study we decided on a self-tracking application, which tracks the user's mood at regular intervals via notifications. This decision is based on the fact that a human's mood can change frequently over a day, which justifies short intervals between notifications. Basically, any notification-triggering application could be used for this research; actual mood assessments are not relevant for this paper and will not be presented or discussed.

The application is called MoodDiary, runs on Android, and is separated into two parts. The first part is a background service that regularly triggers notifications and asks the users to assess their current mood. The second part is an activity that provides an overview of all mood assessments and thereby creates an actual value for the user. We paid attention that the application is of good quality, stable, and reliable. We created an appealing logo and application description to attract many users. The application description makes the user aware that it collects anonymous sensor information for study purposes.

Trigger New Notifications

The background service triggers a new mood assessment notification at regular intervals of about three hours and fifteen minutes. The regular sampling strategy is realistic for such an application and eventually allows us to cover a whole day. The service runs continuously and automatically restarts on phone reboots. For study purposes the service invisibly infers the user's context 15 seconds *before* the actual notification is triggered (see Figure 1, Step 1). After context gathering, a snapshot of the context is taken and the actual notification is shown, which basically looks and behaves similar to, e.g., traditional SMS notifications (see Figure 1, Step 2 and Step 3). The notification takes the phone profile into account, i.e., no audio alert is played if the phone is muted. A user can either answer a notification or reject it. After one minute the notification is automatically dismissed and will no longer be shown in the notification bar.

If a user opens a notification, a sequence of two dialogues comes up (see Figure 1, Step 4 and Step 5). Those dialogues consist of a statement at the top and a 5-point Likert scale in the middle, which ranges from *strongly agree* to *strongly disagree*. The first dialogue gives a statement about a certain mood. The supported mood types are based on findings of Charles Darwin and Paul Ekman: angriness, fearfulness, happiness, sadness, how disgusted and how surprised someone feels. Thus, a possible first statement could be “I feel surprised”. The second statement assesses the inquiry's obtrusiveness: “The previous inquiry was obtrusive”. After the user clicks the OK button in the second dialogue, the dialogue disappears and the assessment is stored in a local database.

Independently of whether a notification was attended or not, we log information to our servers (see Figure 1, Step 6). We log the inferred context information (as described earlier), if a notification has been answered, the mood type, the given answer, and the rated obtrusiveness of each inquiry. We further log information like the device type, the set locale etc. All information is collected anonymously with the users consent.

Value for the User: Overview of Mood Assessments

Beside the background service, the application also provides an activity that shows earlier assessments. This activity is decoupled from the background service and consists of two tabs. The first tab contains a list of all answered or rejected inquiries. Each list element shows which mood type was asked for, if this inquiry was answered, how the user self-assessed his mood and how obtrusive this inquiry was rated. The second tab provides the user with graphs how the individual mood types have changed over time. However, this part of the application is not of relevance for this paper.

User Study

We released the MoodDiary application, to the Google Play store in March 2012 and did our analysis six month later. In the following, we investigate the data from two complementary perspectives. First, we present the most relevant descriptive statistics in which contexts a notification is typically attended or ignored. Second, we create a model to predict opportune moments and investigate its performance.

Results: Statistics

The application was installed by 314 users over a period of five month. Overall 15926 issued notifications were recorded by our server. To exclude users that installed the application without really using it, we applied a filter. We excluded all users that used the application less than 1 day and had an answering rate, i.e., ratio between answered and issued notifications, of less than 10 %. After filtering, a set of 79 users and 6581 issued notifications remained.

About two thirds of the 79 contributing users had an American time zone (GMT/-4 to GMT/-9). Further, 73 users (92.40 %) had set an English locale. Of the 6581 total notifications 1508 were answered and 5073 were not answered, which results into an overall answering rate of 22.91 %. On average, each user was confronted with 83.3 notifications (SD 108.30, median 36.0), of which 19.09 (SD 23.09, median 8.0) were answered. These results indicate that almost every fourth notification is answered. The application was used between 1.01 and 76.63 days, which is about 11.01 days (SD 14.87, median 4.74) on average.

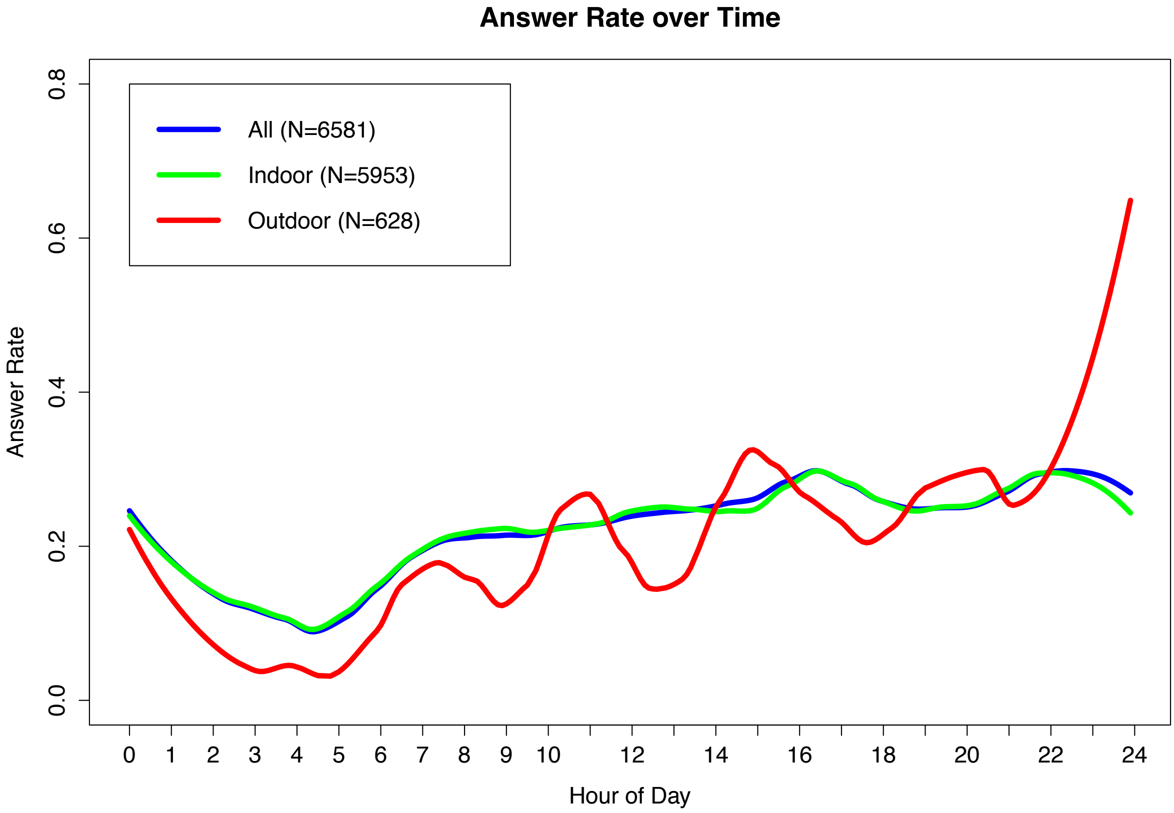

Figure 2: Despite a major drop to 8 % at night times the answer rate is more or less increasing constantly, ending in a global maximum of about 30 % in the evening hours. For outdoor notifications we observed some local minima, which can most likely be credited to commuting.

Time of Day

We were able to collect equally distributed data for the time of day (skewness s=0.42, excess kurtosis k=-0.72). In the following, we investigate if the time of day has an effect on the notification answering behavior. We found that only 0.08 % of the notifications are answered at around 04:24, and 0.31 % are answered at 22:18 (see Figure 2). In relation to all sampled notifications this answering behavior is skewed to the right, i.e., towards the evening (s=-0.61, k=-0.68). The mean time for answered notifications is 14:16 (SD 06:36, median 15:06), the time for unanswered notifications is 12:12 (SD 07:00, median 12:22). A Welch-adapted two-tailed Student t-test indicates that this difference is significant ($p<0.01$). We interpret these statistics in a way that people mostly attend notifications in the later hours of the day.

Location Provider

The location provider can give insights whether the user is situated indoors or outdoors, as a GPS location is most likely to be inferred only in outdoor situations. Consequently, we split our data into an *indoor* data set (no location information and network location) and an *outdoor* data set (GPS locations only). We observed 5953 (90.46 %) indoor and 628 outdoor notifications. For both settings the distribution of observed notifications looked similar, i.e., less notifications at night (see Figure 2). However, for outdoor notifications we identified noticeable local minima throughout the day, which probably can be credited to commutes (09:00 and 17:30) or going for lunch (13:00). 1352 (22.71 %) of the indoor notifications and 156 (24.84 %) of the outdoor notifications were answered. A Chi-squared test indicated that there is no significant difference ($\chi^2=1.12$, $p=0.29$). Thus, we cannot argue that outdoor notifications were answered more likely than indoor notifications or vice versa.

Position Accuracy

Similar to location, the position accuracy can indicate whether a user was located outdoors or indoors when attending or ignoring a notification. The average position accuracy for answered notification is 851.02 m (SD 1174.14 m, median 96 m), whereby it is 1017.41 m (SD 1219.38 m}, median 672 m) for ignored notifications. This difference is significant ($p<0.01$).

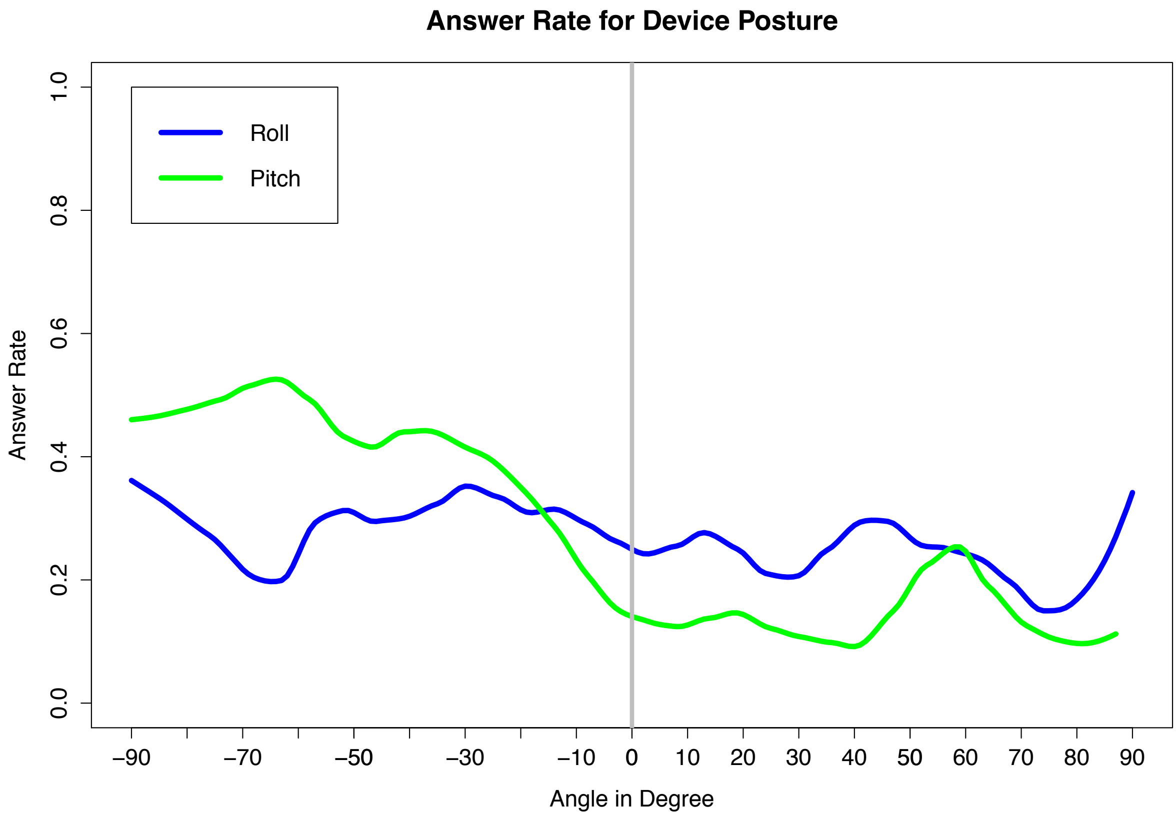

Figure 3: About every second notification is answered if the device is pitched towards the user. 50 % of all answered notifications have a pitch angle between -0.50° and -55.93°.

Device Posture

Roll and pitch angles, which we both recorded for each notification, can be used to estimate the device posture. Roll represents the left/right tilt of the phone, and pitch indicates the up/down tilt. The posture could give indications on how the phone was located, e.g., lying flat on a table, when a notification was answered. The average pitch angle for answered notifications was -31.82° (SD 53.58°, median -28°), while it was -10.89° (SD 67.51°, median -0.39°) for ignored notifications. A t-test showed that this difference is significant ($p<0.01$). The roll angle for answered notifications was 0.31° (SD 22.42°, median 0°), for unanswered 2.20° (SD 22.28°, median 0.02°). This difference is also significant ($p<0.01$).

A visualization of the answer rate over roll and pitch (see Figure 3) showed that the answer rate is about 26.46 % for almost any roll angle. However, for pitch a major peak at around -64° could be observed, where 52.58 % of all notifications were answered. That means that more than every second notification is answered if the device is tilted towards the user by about 60 %. Interestingly, these pitch angles are representatives for how a device is typically held in a user's hand [13, 14].

Proximity

The proximity sensor indicates whether the display is covered by something or not, which is typically used to avoid unintended touch interactions, e.g., during a phone call. However, it can also be used to understand where a phone is probably located, e.g., the sensor is covered in a pocket and not covered in the users hand. 1332 (88.26 %) of the answered and 3816 (75.22 %) of the unanswered notifications came with the information that the display was not covered. A Chi-square test indicated that there is a relation between the facts that a notification was answered and that the proximity sensor was covered (χ^2=116.48, p<0.01). Thus, the display was more likely to not be covered for answered notifications.

Heading, Speed, and Light Level

The headings for answered notifications were 54.46° (GPS, SD 102.60°) and 176.89° (compass, SD 111.82°) and for ignored notifications 55.82° (GPS, SD 98.42°) and 170.65° (compass, SD 107.67°). The speed had a notable variation, but a median of 0.00 m/s for answered and unattended notifications. The median light level for answered notification was 4.09 lux and for ignored notifications 4.00 lux. None of these differences were significant, so we cannot conclude anything.

Data Mining

Beside the statistical investigation of individual measures we processed the obtained data in Weka, a tool for data mining. We wanted to create a classifier, which ideally will be able to predict, based on sensor data, if it is likely that a user will answer an inquiry or not.

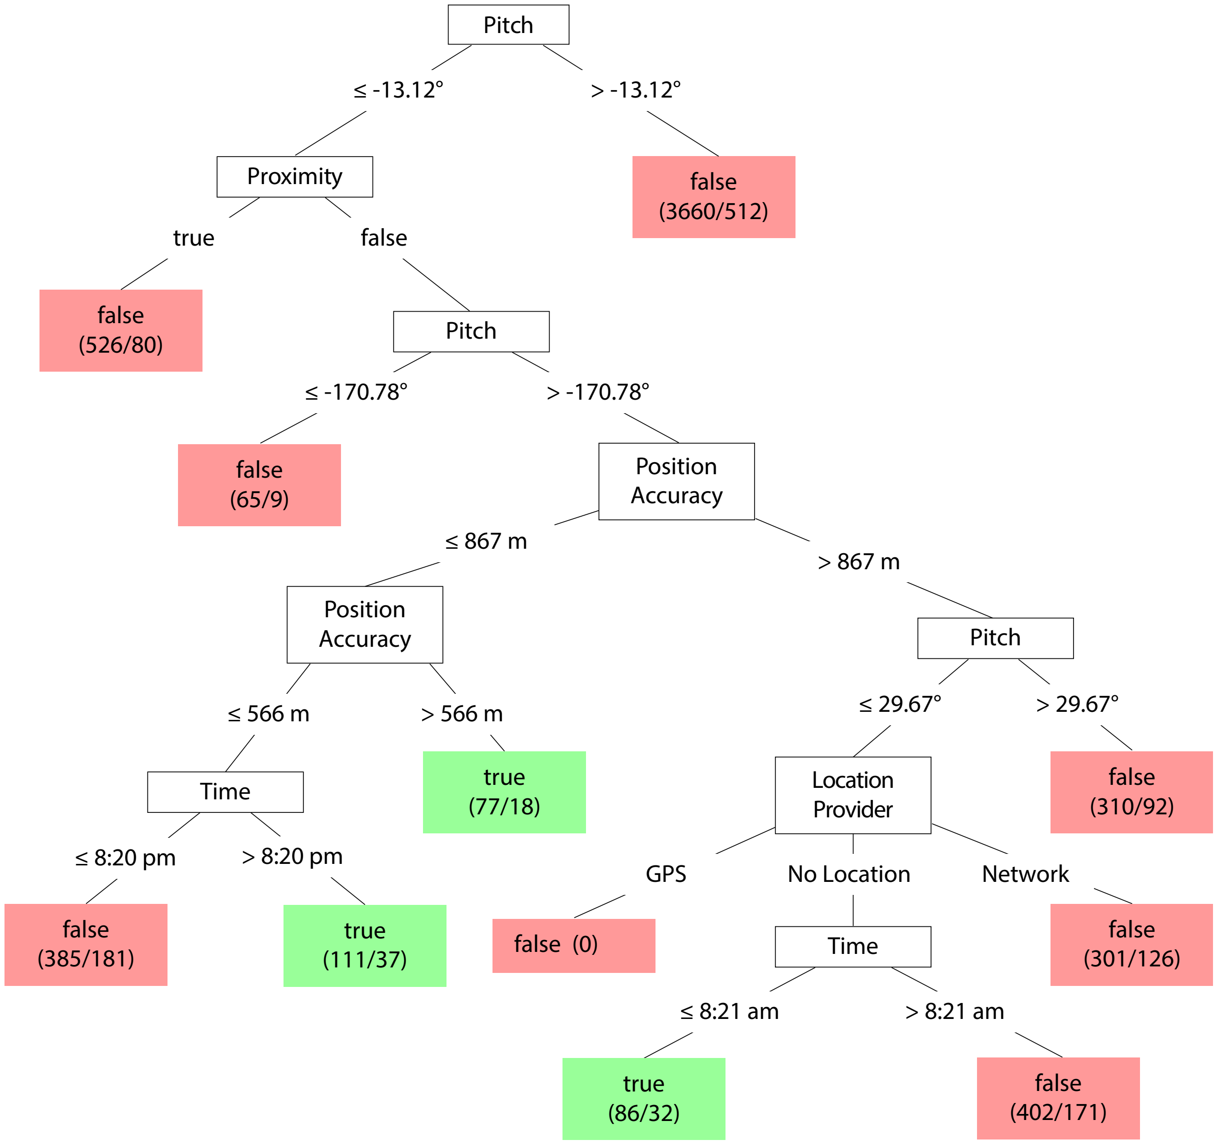

Figure 4: The resulting C4.5 decision tree consists of 20 elements and 11 leaves. It allows to classify whether a notification should be issued (true) or not (false) with an accuracy of 77.85 %.

We investigated different classification algorithms, which we provided with nine attributes to build the model, i.e., location provider, position accuracy, speed, roll, pitch, proximity, time, light level, and whether the request was answered. The full set of 6581 notifications was used to train the classifiers and we used a ten-fold cross-validation procedure for evaluation. We used a trivial classifier as baseline, which predicted that each element belongs to the largest class in the sample, i.e., *false*, and thereby reached an accuracy of 77.08 %.

Our empirical investigation showed that the tree-based C4.5 classifier (see Figure 4) performed best and classified 5114 of the given 6581 notifications (77.85 %) correctly. The precision of the classifier was 0.74 and recall was 0.78. We got 257 true positives, 207 false positives, 1251 false negatives, and 4866 true negatives. This showed that particularly situations where the user would actually answer a notification were misclassified.

Although the accuracies of the baseline classifier and the C4.5 classifier look quite similar, the C4.5 classifier has the overall better performance. While the baseline classifier is unable to predict a single opportune moment, the C4.5 classifier predicts

257 opportune moments accurately. The calculated Kappa statistic of 0.17 and the receiver operating characteristic, area under curve (ROC AUC), of 0.72 indicate that this classifier is not performing outstandingly well, but is definitely better than the baseline condition.

Alternative classifiers didn't perform better: a JRip classifier reached 77.48 % (precision 0.74, recall 0.77), and a neural network reached 76.84 % (precision 0.73, recall 0.77) in a similar evaluation procedure. The overall accuracy of all classifiers is comparable to existing classifiers for office settings [3].

Discussion

In the following we discuss how these results should be interpreted and what are the practical and theoretical limitations of the identified model. Again, we want to emphasize that our set of sensors is unable to measure all aspects that make a mobile situation. Consequently, our findings are only valid for this simplified view on the context.

Trigger Notifications at the Right Time

In the statistical analysis we could observe that the average time of day for answered notifications, 14:16, is significantly later than for ignored notifications, 12:12. Further, the answering rate decreases significantly during nighttime. Two leaves of the decision tree, which classify a notification likely to be answered, have a time node as predecessor. Thus, time seems to be a good predictor.

A combination of both, statistics and decision tree, implies that the ideal time to trigger a notification is before 08:21, after 20:20, but not during nighttime. These times appear logical if we compare them with a typical 9-to-5 working day. Notifications are likely to be answered after getting up, during breakfast or during the commute in the morning. Further, they are likely to be answered when the user is back at home, and the working day is over.

Trigger Notifications When Users Hold the Phone in Their Hands

The identified average pitch angle for answered notifications indicates that the users often already held the phone in their hands when answering a notification. Further, the proximity sensor is significantly less often covered for answered notifications. The derived decision tree is able to reach a high accuracy with pitch and proximity as two important identifying measures. Particularly the fact that 3660 (55.61 %) notifications with a pitch angle greater than -13.12° are classified as unlikely to be answered is particularly remarkable (see Figure 4). We interpret these measures, i.e., pitch and proximity, as an indicator that notifications are more likely to be answered if the users already hold the phone in their hands. This is the most relevant aspect to predict whether a notification will be answered or not.

Relevance and Applicability of the Model in Practice

From a technical perspective, the derived model can be transferred into simple comparisons, which can be executed on today's mobile phones. This makes an application feasible and probable. The presented model gives a basic recommendation in which situations a notification is likely to be answered and in which not. As explained, the approach is limited by the fact that we just measure and include a tiny fragment of the overall context. The identified performance measures indicate that the predictions of the simplified model are only a medium improvement over random predictions. That means that still a considerable number of notifications will probably occur in inopportune moments. Consequently, we think it will be beneficial to further improve the model before application.

We think that our model can serve as a starting point for less simplified models, which incorporate more complex relationships and constraints between individual predictors. These advanced models will probably lead to an increased prediction accuracy and, eventually, to a generalizability among use-cases and an applicability of the approach in everyday life. In addition, given how different a daily routine can be among users, we think it would be beneficial to have user-specific models.

Further, also notifications have different characteristics that can be considered. One example is the relevance of a notification for the user. For particularly relevant, important or time-critical notifications the model, as presented here, is hardly applicable. To make it applicable, much more information about the users and their tasks need to be collected and incorporated. Further, the model needs to be integrated into a complex notification management framework, like sketched by Iqbal et al. [15], that is responsible to issue a notification at opportune moments, but follows a complex set of dynamic constraints.

We think that the resulting model can also be applied and used beyond mediated interruption techniques. For example, we can think about a negotiated interruption technique, where the user's choice if a notification should be handled right away or later can be used to re-fine the model at runtime. This approach would allow the model to adapt to changing user contexts, e.g., if the user starts a new job or moves to a new city and then follows other daily routines.

Limitations

One limitation for large-scale studies via mobile applications is the data validity. Although we did sanity checks, erroneous data samples might still affect the analysis. Further, we cannot say anything about the users' motivation to answer or ignore/reject notifications. Thus, it could be the case that a user ignores a notification although it was issued in an actual opportune moment, thereby influencing our findings. However, we argue that this is regularly not the case and for the majority of the recorded data the user is behaving as initially expected.

Conclusion

In this work we have studied in which contexts a mobile notification is attended and in which it is rejected or ignored. We did so by publishing the mood-tracking application MoodDiary in the Google Play store, where we collected 6581 notifications from 79 users over periods of up to 76 days. This approach gives our studies the advantage that they are done in a real, longitudinal, and large-scale setting, which is beneficial for the results' validity and which is something that has not been done before in interruption research.

We analyzed the notifications from a statistical perspective and did a data mining analysis. We identified that notifications should be triggered at the right time and when the device is already at hand. We illustrate that our derived classifier can be applied already, but a surrounding notification management framework and less simplified models would be helpful to achieve a practical value in day-to-day use. We envision that the applied model can lead to a substantial decrease of notifications in inopportune moments, which therefore reduces the number of annoying and unpleasant interruptions at no costs. This eventually leads to fewer context switches and therefore to less stress, frustration, time pressure and effort [1].

In our future research we want to investigate the actual application of the derived generic model and compare it against user-specific models. Further, we plan to incorporate non-sensor information, like phone activities or calendar entries, in future model revisions. Eventually, we want to embed the model in a holistic notification management framework to investigate options how urgent notifications can be treated. All future work will be done to pursue the objective to further improve the user experience.

Acknowledgements

The authors are grateful to the European Commission and the German Federal Ministry of Education and Research, which co-fund the AAL JP project NavMem (AAL-2011-4-122).

References

[1] Gloria Mark, Daniela Gudith, Ulrich Klocke, The cost of interrupted work: more speed and stress. Proceedings of CHI 2008.

[2] James Fogarty, Scott E. Hudson, Jennifer Lai, Examining the robustness of sensor-based statistical models of human interruptibility. Proceedings of CHI 2004.

[3] James Fogarty, Scott E. Hudson, Christopher G. Atkeson, Daniel Avrahami, Jodi Forlizzi, Sara Kiesler, Johnny C. Lee, Jie Yang, Predicting human interruptibility with sensors. ACM Trans. Comput.-Hum. Interact. 12(1) 2005.

[4] Yoshiro Miyata, Donald A. Norman, Psychological Issues in Support of Multiple Activities. User Centered System Design: New Perspectives on Human-Computer Interaction, 1986.

[5] Daniel Avrahami, James Fogarty, Scott E. Hudson, Biases in human estimation of interruptibility: effects and implications for practice. Proceedings of CHI 2007.

[6] Daniel C. McFarlane, Kara A. Latorella, The scope and importance of human interruption in human-computer interaction design. Hum.-Comput. Interact. 17(1) 2002.

[7] Daniel C. McFarlane, Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. Hum.-Comput. Interact. 17(1) 2002.

[8] Piotr D. Adamczyk Brian P. Bailey, If not now, when?: the effects of interruption at different moments within task execution. Proceedings of CHI 2004.

[9] Brian P. Bailey, Shamsi T. Iqbal, Understanding changes in mental workload during execution of goal-directed tasks and its application for interruption management. ACM Trans. Comput.-Hum. Interact. 14(4) 2008.

[10] Joel E. Fischer, Chris Greenhalgh, Steve Benford, Investigating episodes of mobile phone activity as indicators of opportune moments to deliver notifications. Proceedings of MobileHCI 2011.

[11] Joyce Ho, Stephen S. Intille, Using context-aware computing to reduce the perceived burden of interruptions from mobile devices. Proceedings of CHI 2005.

[12] Nicky Kern, Bernt Schiele, Context-Aware Notification for Wearable Computing. Proceedings of ISWC 2003.
[13] Ken Hinckley, Jeff Pierce, Mike Sinclair, Eric Horvitz, Sensing techniques for mobile interaction. Proceedings of UIST 2000.
[14] Beverly L. Harrison, Kenneth P. Fishkin, Anuj Gujar, Carlos Mochon, Roy Want, Squeeze me, hold me, tilt me! An exploration of manipulative user interfaces. Proceedings of CHI 1998.
[15] Shamsi T. Iqbal, Brian P. Bailey, Effects of intelligent notification management on users and their tasks. Proceedings of CHI 2008.

Short Author Bios

Benjamin Poppinga is a researcher in the Interactive Systems Group at OFFIS - Institute for Information Technology. At the same time he is a doctoral student at the University of Oldenburg, Germany. Benjamin's research focuses on mobile "in the wild" context sensing and observation techniques, especially in a mid- to long-term perspective.

Dr. Wilko Heuten is a researcher at OFFIS - Institute for Information Technology since 2001. Since 2008 Wilko Heuten leads the group Interactive Systems in the health department at OFFIS. His main research is targeted to the design, development and evaluation of new interaction methods of traditional and beyond traditional desktop applications.

Prof. Dr. Susanne Boll is Professor for Media Informatics and Multimedia Systems in the Department of Computing Science at the University of Oldenburg, Germany, and she is a member of the board of the OFFIS - Institute for Information Technology. Her research interests lie in the field of multimodal multimedia retrieval, context-aware mobile multimedia systems and multimodal multimedia user interfaces.

Complete Contact Information

Benjamin Poppinga
OFFIS – Institute for Information Technology
Escherweg 2, 26121 Oldenburg, GERMANY
benjamin.poppinga@offis.de

Wilko Heuten
OFFIS – Institute for Information Technology
Escherweg 2, 26121 Oldenburg, GERMANY
wilko.heuten@offis.de

Susanne Boll
University of Oldenburg, Media Informatics and Multimedia Systems Group
Escherweg 2, 26121 Oldenburg, GERMANY
boll@informatik.uni-oldenburg.de