Ubiquitous and Mobile Computing CS 528: *Designing Content-driven Intelligent Notification Mechanisms for Mobile Applications*

Adam Chaulk
Zhongyuan Fu

Computer Science Dept.
Worcester Polytechnic Institute (WPI)

Outline

- Introduction
- Related Work
- Methodology
- Evaluation/Results
- Challenges
- Future Work
- Conclusions

Introduction

- Notifications can be annoying!
- What if?
- Objective: find the most opportune moment to deliver notifications

Motivation - What is an Opportune Moment?

- Opportune only when the user will answer it immediately
- Aim to reduce response time of the user and the acceptance rate of notifications

The Solution - NotifyMe

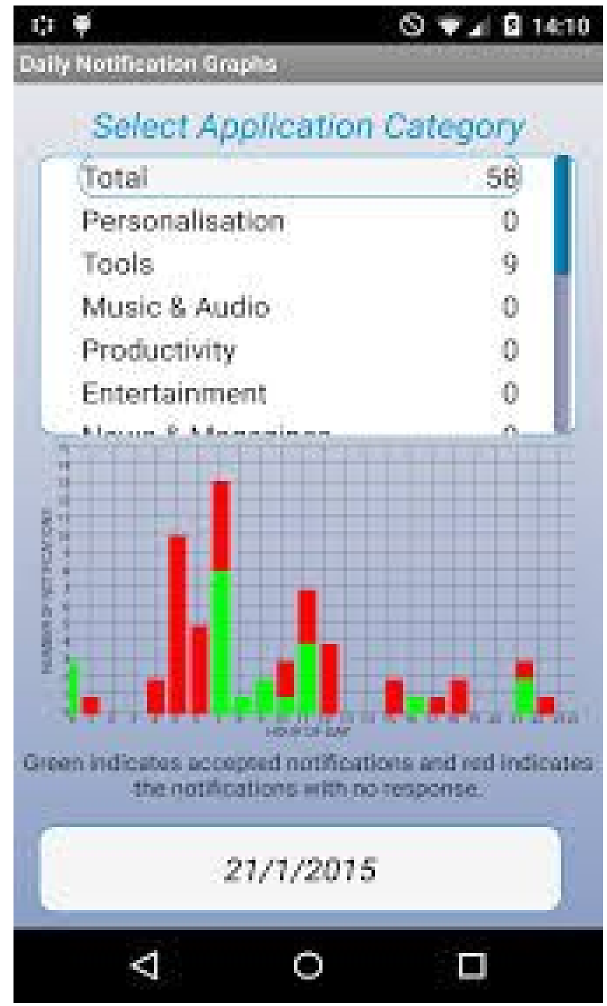

Related Work

- Similar Studies
  - Contextual factors used to infer interruptibility with
    - Contextual transitions [Ho, et al]
    - Engagements on device [Fischer, et al.]
    - Time of day, location, and activity [Pejovic, et al]
- How is this different?

The Study

- In-the-wild notifications
- 35 users
  - Published on Google Play Store
  - Ages 21-31
  - Advertised at University of Birmingham (UK)
- 3 weeks, 70,000 notifications, 4,096 questionnaire responses

Question	Options
How would you rate the notification content?	Likert scale rating between 1 and 5 (1 = very annoying and 5 = very interesting).
Where would you like to receive notifications with similar content?	Home, workplace, other, anywhere and I don't want.
When would you like to receive notifications with similar content?	Morning, afternoon, evening, night, anytime and never.
How are you feeling?	Happy, sad, bored and annoyed.
Are you busy?	Yes and no.
Where are you?	Home, workplace, public, other.

Table 2. Questions and their options from NotifyMe questionnaire.

Methodology

- Data collection forms:
  - Measures notification responses (accept/decline)
  - NotifyMe notifications
  - Questionnaires
- Google’s Notification Listener Service to trace notifications, and Activity Recognition API, ESSensorManager to get context info

Dataset

- Classified by info type
  - Work
  - Social
  - Family
  - Other
- Generated label-to-notification map

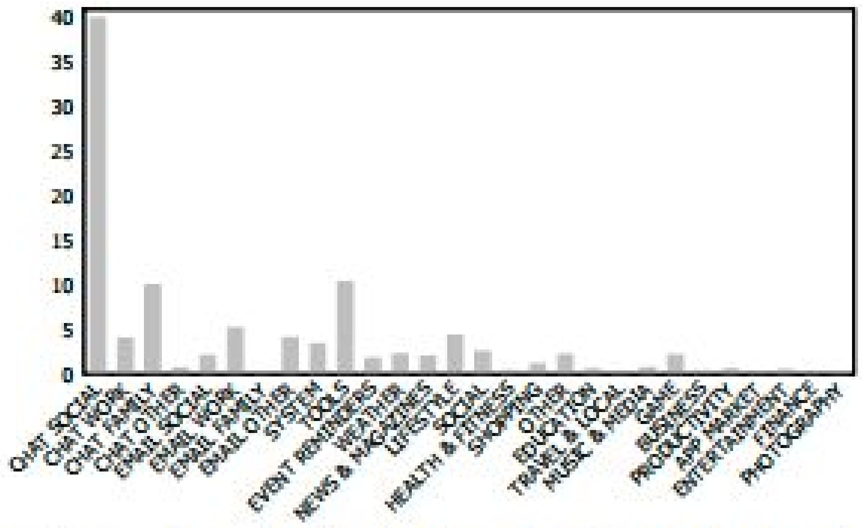

Figure 2. Percentage of notifications for each category and sub-category. The sub-categories are derived by using the recipient's relationship with the sender.

Design Tradeoffs

- Tradeoffs
  - Privacy concerns - only access message titles
  - Notification categorization
  - User opt-out
  - “Accepting” a notifications = launching the app

Results

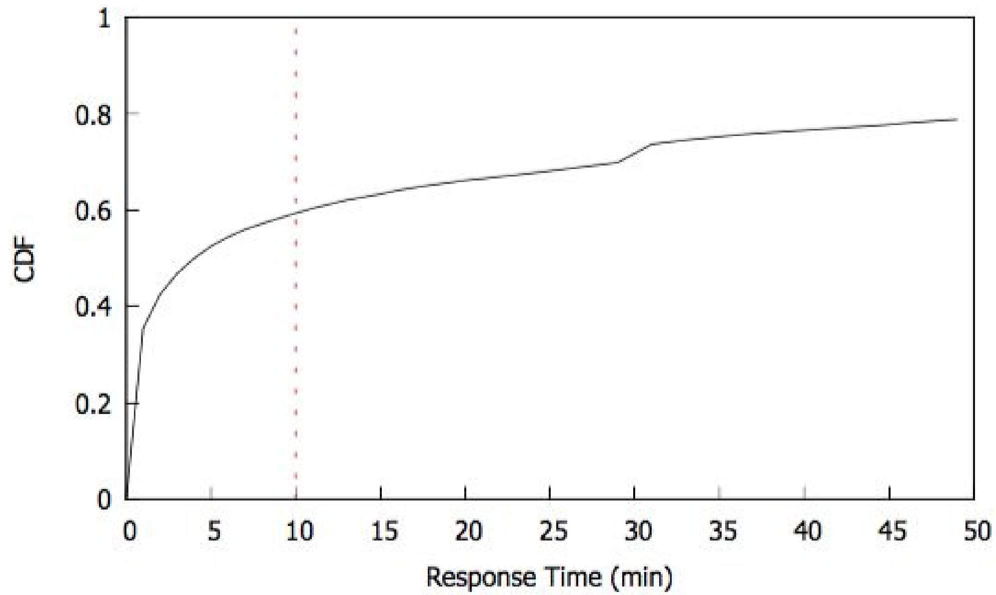

Figure 3. CDF of response time for notifications.

- Collected 70,000 notification samples
- More than 60% notifications were clicked within 10 minutes from the time of arrival

Impact of Context on Response Time

- Location
  - home, workplace, the other
- Surrounding sound
  - silent or speaking
- Activity
  - still, on foot, on bicycle, in vehicle

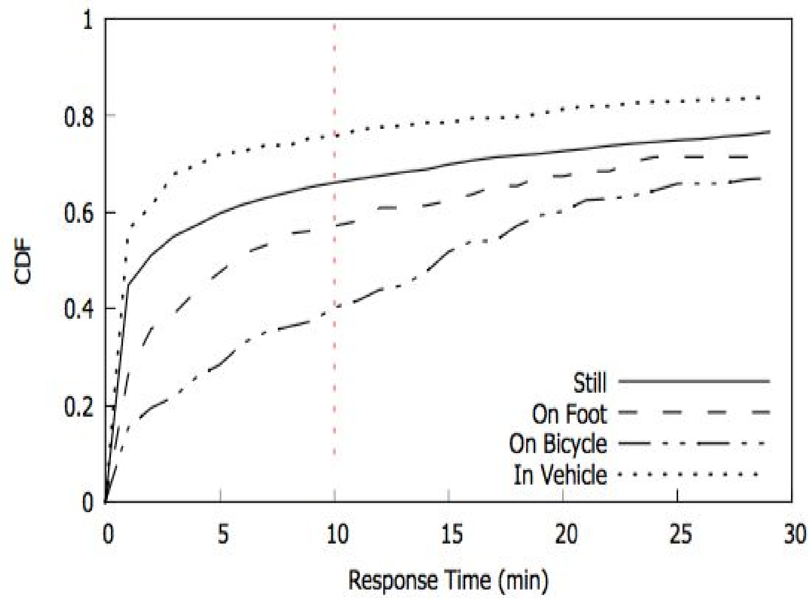

Figure 4. CDF of response time for notifications received while performing different activities.

Impact of Content on Notification Acceptance

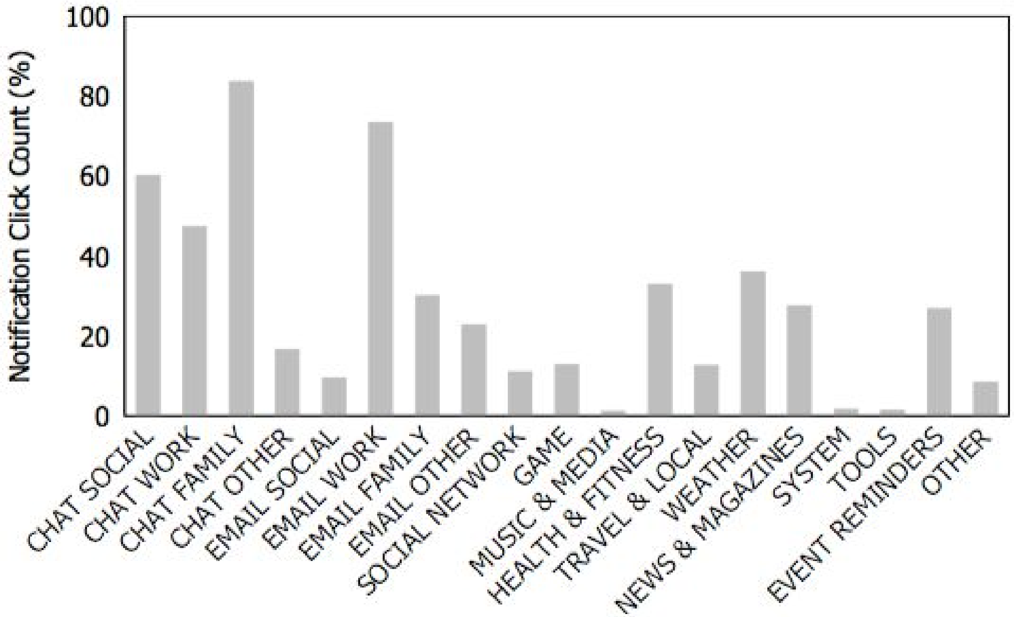

Figure 5. Click count percentages for the notifications of each category.

- The notification from different categories have a varying acceptance rate
- E.g. Chat Family, System, Tools, Music & Media

Building the Prediction Model

- Build models for predicting notification acceptance rate by three algorithms: Naive Bayes, AdaBoost, and Random Forest

- Two approaches for building prediction models
  - Data-driven learning
  - User-defined rules

Approaches for building the Prediction Model

- Data-driven learning that relies on the evidences rather than personal intuition
  - without using information type and social circle
  - using only information type
  - using information type and social circle

- User-defined rules that rely on the user's own intuitions
  - notification category
  - best location
  - best time

Evaluation

- Sensitivity
  - # of predicted accepted notifications / total # of accepted notifications
- Specificity
  - # of predicted declined notifications / total # of declined notifications

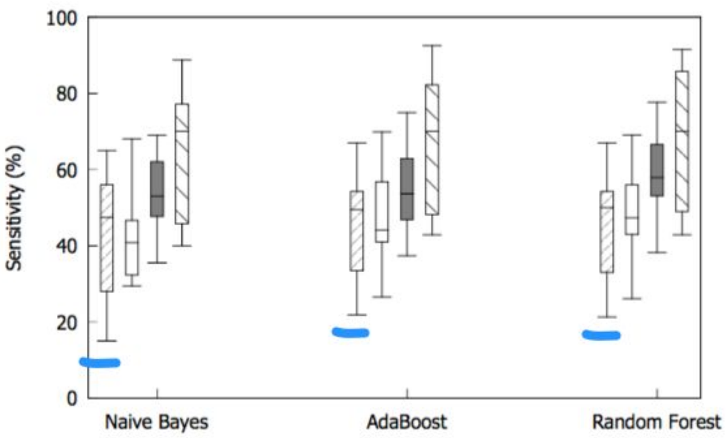

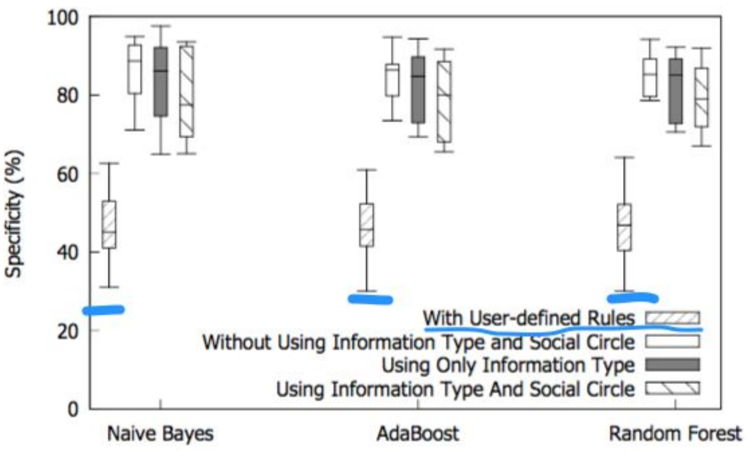

Prediction results of the predictors trained by using 3 different set of features for data-driven learning and user-defined rules

Generic vs Personal Behavioral Model

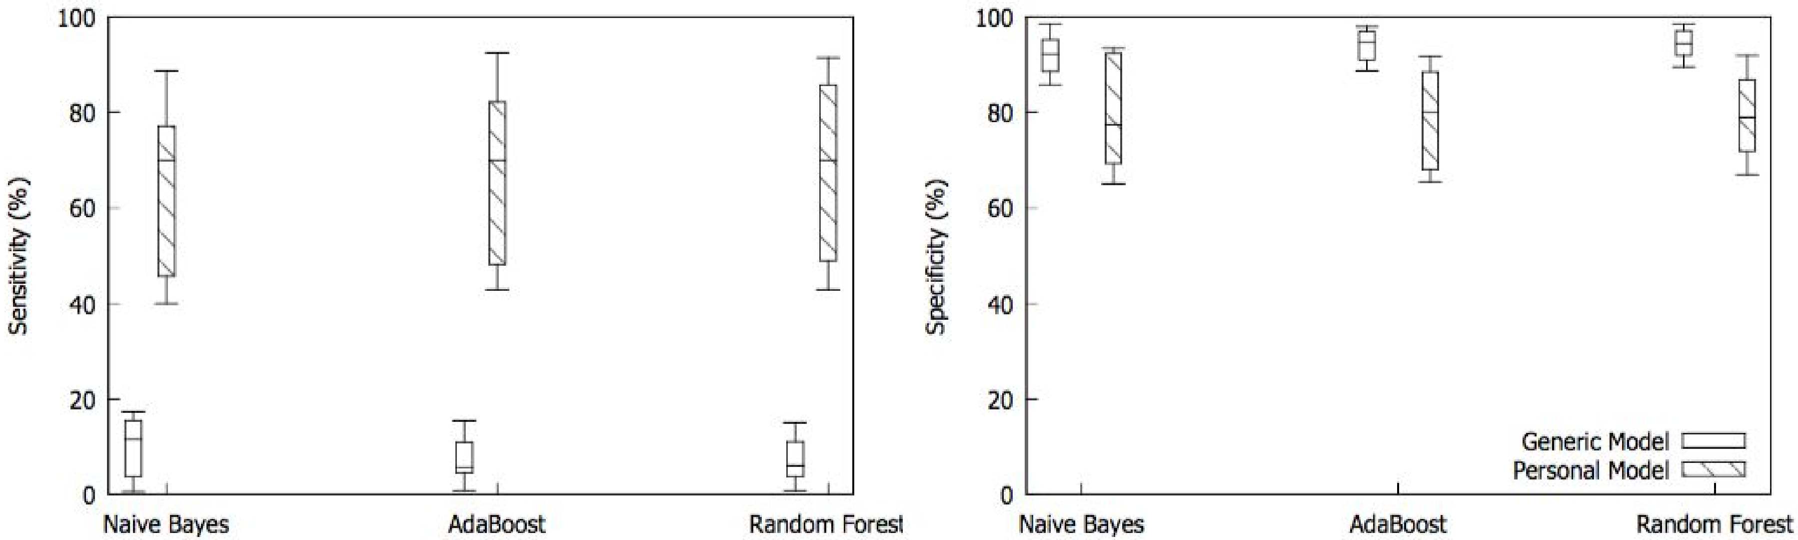

Figure 7. Prediction results of the generic and individual-based models.

Limitations

- Notification might be attended on another device
- Some applications do not require any action taken, just dismiss

Challenges

- Multiple notifications from one application at once

Conclusions

- A user's activity can impact the time delay in the response to a notification
- The chat notification, where the sender is a family member or a relative of the user, have the highest acceptance rate
- The acceptance value of notifications vary for each category
- The acceptance of a notification within 10 minutes from its arrival time can be predicted with an average sensitivity of 70% and a specificity of 80%

Future Work

- Increase accuracy of the prediction model
  - Better-defined categories
  - Use Natural Language Processing
- Privacy

References

- *http://www.ucl.ac.uk/~ucfamus/papers/ubicomp15_designing.pdf*
- *http://dl.acm.org/citation.cfm?id=2037402*
- *http://web.media.mit.edu/~intille/papers-files/HoIntille05.pdf*
- *http://www.ucl.ac.uk/~ucfamus/papers/ubicomp14.pdf*

Questions?