Improving the Mobile Phone Habitat - Learning Changes in User's Profiles

Robert Bridle and Eric McCreath

Department of Computer Science, Faculty of Engineering and Information Technology, The Australian National University

Abstract. Mobile phones are becoming a popular platform for a range of applications. However, due to size restrictions, the interfaces of these applications can be difficult to use. Customising an interface for a particular user offers the potential to improve an interface's efficiency. In this paper, we propose customising a mobile phone's Profile application. We apply a machine learning approach to discover concepts that describe a user's profile-activations in terms of their scheduled appointments. We found that it is possible to learn useful concepts, which maybe used to improve the users interaction with mobile phone devices.

1 Introduction

More Australians have access to mobile phones than desktop computers[1]. However, as noted by Amant, Horton and Ritter[6], user interfaces on such devices are often difficult to use. In this paper we focus on a small but important application, the Profiles application. The profile of a mobile includes attributes such as: ringing tones, ringing volume, vibration on/off, keypad tones, and SMS alerts. These profiles enable users to quickly change the setting of the phone depending upon the context they find themselves in. Yet it is easy to forget to change the profile of the phone for the current context. e.g. your phone rings in the middle of a movie, or you miss a number of important calls because you left the phone in *Silent* mode. In this paper we investigate how a direct application of a standard Bayesian learning approach may be used to learn when and how a user changes their profile over time. This domain provides a number of interesting challenges for machine learning because the training data involves temporal attributes, and the training set size is very small. These challenges have been touched on briefly in previous work, most notable that of on-line command-line prediction[2,4] and in Mitchell's calendar apprentice system[5].

2 The Profiles Application and Learning Approach

The Profiles application exists on nearly all mobile phones. The Profile application consists of a set of profiles. Each profile defines properties including: ring tone, ring volume, SMS alert tone, keypad tones and warning tones. The Profiles application is used to change which profile is active. The Profiles application we

S. Zhang and R. Jarvis (Eds.): AI 2005, LNAI 3809, pp. 970–974, 2005.

investigated included the default profiles: *General*, *Meeting*, *Silent*, *Outdoors*, *Offline* and *Pager*. Each profile defines properties consistent with the activity that their name suggests. Other user-defined profiles may be added. A problem with the Profiles application is that a user has to select which profile they want active every time their situation changes. However, often the profile required by a user is dependent upon the activity they are undertaking. We propose a learning approach that can determine the relationship between a user's activities and the profile they want active. Automating profile-activations would save the user from specifying rules and reduce interactions with their mobile phone.

We investigated using a mobile phone's *Calendar* application, to gather information on a user's scheduled appointments. We considered determining a concept that would describe the relationship between a user's appointments and the profile they activated. An obvious example of a user's schedule determining the profile they activate, would occur when a user schedules a meeting. It is likely a user will select the *Silent* profile for the duration of the scheduled meeting and the *General* profile afterwards. In addition, it is unlikely that a user will activate the *Silent* profile exactly on an appointment's scheduled start-time or select the *General* profile exactly on an appointment's scheduled finish-time. Any learning approach must not only determine which appointments govern which profile-activations, but also the amount of time appropriate between when a profile activates and when an appointment starts or finishes.

We implemented a profile monitoring application on a Nokia™Series 60 mobile phone. Whenever a user activated a profile a training example was generated. Each training example was represented by a set of attributes and a class label. The attributes consisted of: the time, the location (nearest cell-tower) and the details (subject and location) of appointments scheduled for that day. The class label represented the profile that was activated, either: *General*, *Meeting*, *Silent*, *Outdoors*, *Offline* or *Pager*. A probabilistic approach was used to generate a concept for each class[3]. From all training examples gathered over a week, a user's week-long calendar schedule was constructed. For each class, we considered every possible model of profile-activations about the user's schedule. We also considered minute time-periods, that ranged from 20 minutes before to 20 minutes after appointment start or finish times. The most likely model became the target concept. We assumed that the likelihood of each of model was given by:

$$L(model) = \lambda^k (1 - \lambda)^{n-k} \tag{1}$$

where n is the number of times the model could correspond to a training instance, and k is the number of times the model does correspond to a training instance. The value λ represents the probability that the profile activation is carried out for the appropriate scheduled item. For each model, and for all possible values of λ, we determine the model and λ values of highest likelihood. For each class, we present the user with the most likely models for each appointment subject. These models are the induced hypotheses.

The probabilistic approach described has a number of properties that make it suitable for the intended learning task. Firstly, we assumed that users would sometimes forget to activate a profile. A probabilistic model makes the learning approach robust to this kind of noise. Secondly, a probabilistic model allows us to quantify the uncertainty in learnt concepts. We can make decision on the usefulness of concepts based on this measure of uncertainty. Finally, probabilistic models can be highly parameterised, enabling us to place a strong bias on the learner's hypothesis language. This allows us to restrict the possible concepts considered by the learner.

3 Experimental Results and Discussion

We gave 4 users a Nokia$^{\text{TM}}$phone. The phone was installed with our profile monitoring application. The profile monitoring application ran unobtrusively in the background and was used to gather training examples (as described in Section 2). We monitored each user's profile-activations over a period of one week. The learning approach described in Section 2 was then invoked off-line on the training examples collected over the week. Table 1, shows the concepts generated for each user. Interpreting these parameters we find that User 1 is expected to activate the *Silent* profile 3 minutes after the start of appointments with the subject [Pathology], 5 minutes after the start of appointments with the subject [Movie] and so on. An important question, is whether the concepts generated would be of future use to the user. We asked each user whether they would accept or reject the generated concepts. For those concepts that were rejected, we also asked the user to indicate the reason for the rejection. Table 2 shows the rejected concepts.

Table 1 and Table 2 show that at least one of the concepts generated for each user is accepted. Given the limited amount of data provided to the learner,

Table 1. The concepts identified for the 4 users after a week of training examples

User 1				User 2			
Profile	Subject	Time	λ	Profile	Subject	Time	λ
Silent(start)	[Pathology]	3.0	0.99	Silent(start)	[Auscc]	-2.0	0.90
Silent(start)	[Movie]	5.0	0.99	Silent(start)	[Gym]	-5.0	0.90
Silent(start)	[TeamLunch]	0.0	0.99	Silent(start)	[Seminar]	-5.0	0.90
Silent(start)	[Pds]	-6.0	0.99	Silent(start)	[Sam farewell]	-2.0	0.90
Silent(start)	[JavaTraining]	-20.0	0.99	General(finish)	[Auscc]	6.0	0.90
General(finish)	[Pathology]	-9.0	0.99	General(finish)	[Gym]	8.0	0.90
General(finish)	[TeamLunch]	0.0	0.99	General(finish)	[Seminar]	2.0	0.90
General(finish)	[Pds]	8.0	0.99				
General(finish)	[TeamMeeting]	7.0	0.99				

User 3				User 4			
Profile	Subject	Time	λ	Profile	Subject	Time	λ
Meeting(start)	[AI]	-10.0	0.80	Silent(start)	[Lunch]	-7.0	0.99
Meeting(start)	[OS]	-10.0	0.80	General(finish)	[Lunch]	11.0	0.99
General(finish)	[AI]	10.0	0.80				

Table 2. The concepts rejected by the 4 users, reason are given below each concept. Note, User 3 and User 4 did not reject any of their concepts.

User 1 - Rejected the Following Concepts.

Profile	Subject Parameters	Time	λ
Silent(start)	[Pathology]	3.0	0.99
Silent(start)	[Movie]	5.0	0.99
Silent(start)	[TeamLunch]	0.0	0.99
General(finish)	[Pathology]	-9.0	0.99
General(finish)	[TeamLunch]	0.0	0.99

Reason: *Silent* profile should occur before and not 3 minutes after the start of [Pathology].
Reason: *Silent* profile should occur before and not 5 minutes after the start of [Movie].
Reason: *Silent* profile should occur before and not 0 minutes after the start of [TeamLunch].
Reason: *General* profile should occur after and not 9 minutes before the finish of [Pathology].
Reason: *General* profile should occur after and not 0 minutes before the finish of [TeamLunch].

User 2 - Rejected the Following Concepts.

Profile	Subject Parameters	Time	λ
Silent(start)	[Sam farewell]	-2.0	0.90

Reason: The appointment [Sam farewell] will not occur again.

this is a promising result. It is important to note that the concepts learnt do not completely capture all of the user's profile-activation intentions. For instance, for User 1, User 2 and User 3 in Table 1, the set of appointment subjects in the *Silent* profile concept is different from the set of appointment subjects in the *General* profile concept. If these concepts formed a rule-base, then the *Silent* profile would activate before some appointment subjects and then remain active until the user intervened. A user can not entirely rely on the generated concepts to describe their intentions. Instead, the generated rules would be of most use augmenting a manually defined rule-base.

4 Conclusion

This paper has investigated a method for learning the concepts that govern profile-activations on a mobile phone. The Profile application learning environment posed several interesting challenges. These included training data with temporal attributes, and a very small training set. Despite this, we demonstrated a learning approach that could find useful concepts.

References

1. Australian Bureau of Statistics. Measures of a knowledge-based economy and society. *Australia, Australia Now*, (Cat. no. 8146.0), 2003.
2. B. Davison and H. Hirsh. Probabilistic online action prediction. In *Intelligent Environments: Papers from the AAAI 1998 Spring Symposium, Technical Report SS-98-02.*, pages 148–154, 1998.

3. R. Duda, P. Hart, and D. Stork. *Pattern Classification*. John Wiley and Sons, 2nd edition, 2001.
4. B. Korvemaker and R. Greiner. Predicting unix command lines: Adjusting to user patterns. In *Proceedings of the Seventeenth National Conference on AI and Twelfth Conference on Innovative Applications of AI*, pages 230–235, 2000.
5. T. Mitchell, R. Caruana, D. Freitag, J. McDermott, and D. Zabowski. Experience with a learning personal assistant. *Communications of the ACM*, 37(7):80–91, 1994.
6. R. St. Amant, T. Horton, and F. Ritter. Model-based evaluation of cell phone menu interaction. In *Proceedings of the Conference on Human Factors in Computing Systems*, pages 343–350, 2004.