Jogger: Models for Context-Sensitive Reminding

Ece Kamar
Microsoft Research
One Microsoft Way
Redmond, WA, 98052
eckamar@microsoft.com

Eric Horvitz
Microsoft Research
One Microsoft Way
Redmond, WA, 98052
horvitz@microsoft.com

ABSTRACT

We describe research on principles of context-sensitive reminding that show promise for serving in systems that work to jog peoples' memories about information that they may forget. The methods center on the construction and use of a set of distinct probabilistic models that predict (1) items that may be forgotten, (2) the expected relevance of the items in a situation, and (3) the cost of interruption associated with alerting about a reminder. We describe the use of this set of models in the Jogger prototype that employs predictions and decision-theoretic optimization to compute the value of reminders about meetings.

Categories and Subject Descriptors

I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent agents*

General Terms

Design, Human Factors

Keywords

Reminder systems, user modeling, decision-theoretic reminding

1. INTRODUCTION

In the course of daily life, people often forget information that would be valuable to them if they had remembered it at the right time. We present a study of methods for context-sensitive reminding that hold promise for effective personal reminder systems. The approach employs a set of probabilistic models learned from labeled data that predict a set of outcomes required for effective reminding. These outcomes include (1) the probability that information will not be remembered, (2) the relevance of the forgotten information in a current or forthcoming setting, and (3) the cost of transmitting the reminder to a user within a current context. We shall review the set of models and describe how we combine them into a working prototype named Jogger. Jogger follows a decision-theoretic approach to distinguish reminders that are beneficial for a user's performance from the ones that are not. We highlight key ideas in the context of reminders about meetings.

Several reminder systems have been proposed in previous work [3, 4, 2, 8]. None of these reminder systems employ a principled methodology for identifying the value, relevance, and timing of a reminder–key ingredients for generating effective reminders. Jogger follows the line of research on using decision-theoretic approaches to manage and filter notifications [6].

A more detailed presentation of the ideas investigated in this work, including an evaluation of the Jogger prototype on real-world calendar data, and the extensions of the prototype that reasons about reminder timing and real-time traffic and location information can be found in [7].

Cite as: Jogger: Models for Context-Sensitive Reminding, Ece Kamar and Eric Horvitz, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. XXX-XXX.

2. EXPECTED VALUE OF A REMINDER

Reminders are useful in helping users to recall tasks that need to be accomplished or providing users with other enabling information (e.g., names of people met before in a social setting). An ideal reminder system should consider both the potential benefit of a reminder and the cost of interruption associated with transmitting the reminder. This section discusses how we compute the cost and benefits of a reminder based on predictions about a user's context.

The utility of a reminder for task m depends on the cognitive state of a user: has the user forgotten all or some information that might be included in a reminder? Jogger considers three mental states with respect to recall of information useful for completing tasks under consideration: (1) F^m represents the state in which a user has forgotten all about m, (2) D^m represents the state in which the user has forgotten or is unsure about a subset of details regarding the task, such as its location, start time (or deadline), and other participants, and (3) R^m represents the state in which the user remembers that task m exists and also remembers all of the details regarding the task. Given evidence E that comprises observations about a user's state, $p(F^m|E)$, $p(D^m|E)$, $p(R^m|E)$ are the probabilities of the user being in states F^m, D^m, R^m respectively. F^m, D^m and R^m are mutually exclusive and collectively exhaustive.

The benefit of a reminder depends on the cognitive state of a user. As an example, if a user completely forgets about a meeting, she will not be able to participate nor contribute to a task. If a user forgets some details about a forthcoming meeting (e.g., the location of a meeting), the utility of the outcome may decrease because of tardy arrival. $U^m_F(E)$ and $U^m_D(E)$ represent user's utilities for receiving a reminder for m in states F^m and D^m respectively.

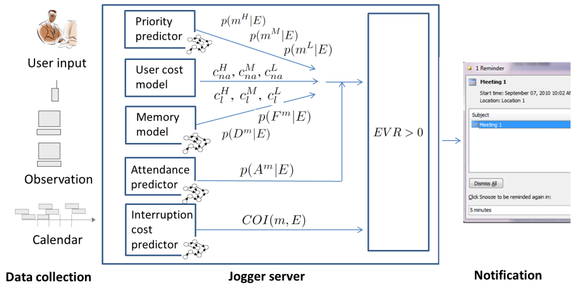

Figure 1: Components of Jogger.

The benefit of a reminder about task m to a user depends on whether m is relevant to the user's plans. $p(A^m|E)$ is the likelihood that the user would engage in task m if she remembers about m. In the meeting reminder context, $p(A^m|E)$ represents the probability of attending meeting m given E, evidence about the meeting. $COI(m,E)$ represents the cost of interrupting the user by delivering a reminder about m, given evidence E about the user's state. We compute the *expected value of reminding* (EVR) as given below:

$$EVR(m) = p(A^m|E)\ (p(F^m|E)\ U_F^m(E) + p(D^m|E)\ U_D^m(E)) - COI(m,E)$$

Next, we formalize $U_F^m(E)$ and $U_D^m(E)$ for the context of meeting reminders. We make the assumption that if a user is in state F^m, the user fails to attend meeting m; if the user is in state D^m, she misses the first t minutes of the meeting because of problems with recalling the details about the meeting; and if the user is in state R^m, the user is on time for the start of a meeting. Jogger system has the priority predictor for inferring for any meeting m the probability that m has high priority $p(m^H|E)$, medium priority $p(m^M|E)$ and low priority $p(m^L|E)$. We ask the user to evaluate the value of time for three possible cases; the minute cost for being late, c_l^H for high, c_l^M for medium, c_l^L for low priority meetings; the total cost for not attending to a meetings, c_{na}^H for a high, c_{na}^M for a medium, c_{na}^L for a low priority meeting, and the minute cost for being early, c.

$$U_F^m(E) = (p(m^H|E)\ c_{na}^H) + (p(m^M|E)\ c_{na}^M) + (p(m^L|E)\ c_{na}^L)$$
$$U_D^m(E) = t((p(m^H|E)\ c_l^H) + (p(m^M|E)\ c_l^M) + (p(m^L|E)\ c_l^L))$$

A schematic view of the Jogger prototype is displayed in Figure 1. Jogger gathers relevant information about a user's context by accessing the user's calendar, by monitoring computer activity, and detecting video and audio signals. The information collected from the data collection component is used for inferences needed to compute the net expected value of reminders. For each reminder opportunity, the system infers the expected value of reminding the user given the inferred cost of interruption, and reminds the user only if the associated value is positive.

3. PREDICTIVE MODELS

Jogger has access to appointments drawn from Microsoft Exchange, along with a constellation of atomic and derived meeting properties that serve as evidential features about the meetings. A set of appointments drawn from several months of an online calendar are composed into a case library of training set of meeting instances. We asked participants to tag meetings with several labels via a tagging tool. Two labels encode a user's assessment about attending a meeting and priority of a meeting. A third label represents whether users would forget about the meeting or about important meeting details. The system generates a training set by combining each meeting instance tagged by a user with a set of attributes acquired from the user's personal Outlook profile. These attributes include the day and time of the meeting, its location and organizer, the response status of the user, and whether the meeting is recurrent.

We perform Bayesian structure learning to build probabilistic models that can be used to predict whether a user has forgotten that a meeting exists, whether a user has forgotten about some details of a meeting, and the relevance and the importance of a meeting [1]. Similar models for predicting meeting importance and relevance have been previously used in the Coordinate system [6].

Jogger uses a two-layer approach to estimate the cost of interrupting a user: activity-based predictions of the cost of interruption inferred by BusyBody [5] and the meeting-based interruptability prediction model of the Coordinate system [6]. By doing so, we can infer the cost of interrupting a user when the user performs office activities, and when the user is in a meeting based on the importance of the meeting.

4. FUTURE WORK

We are exploring several extensions of Jogger, which include (1) deploying the prototype in the open world, (2) improving the predictive models via active learning to focus evidence gathering, and (3) applying the principles of context-sensitive reminding to complex task domains. We believe that the development of personalized reminder systems that come to understand the nuances of users' memories and needs for memory jogging may one day provide great value to people in the course of daily life.

5. REFERENCES

[1] D. Chickering, D. Heckerman, and C. Meek. A Bayesian approach to learning Bayesian networks with local structure. In *UAI*, 1997.

[2] K. Conley and J. Carpenter. Towel: Towards an intelligent to-do list. In *Proceedings of the AAAI Spring Symposium on Interaction Challenges for Artificial Assistants*, 2007.

[3] R. DeVaul, B. Clarkson, and A. Pentland. The Memory Glasses: Towards a wearable context aware, situation-appropriate reminder system. In *Workshop on Situated Interaction in Ubiquitous Computing*, 2000.

[4] A. Dey and G. Abowd. Cybreminder: A context-aware system for supporting reminders. In *Handheld and Ubiquitous Computing*, 2000.

[5] E. Horvitz, P. Koch, and J. Apacible. BusyBody: creating and fielding personalized models of the cost of interruption. In *CSCW*, 2004.

[6] E. Horvitz, P. Koch, C. Kadie, and A. Jacobs. Coordinate: Probabilistic forecasting of presence and availability. In *UAI*, 2002.

[7] E. Kamar and E. Horvitz. Investigation of Principles of Context-Sensitive Reminding. Technical report, MSR-TR-2010-174, Microsoft Research, 2010.

[8] K. Myers and N. Yorke-Smith. Proactive behavior of a personal assistive agent. In *AAMAS*, 2008.