

# Preference Learning for Adaptive Interaction

**Julie S. Weber**

Computer Science and Engineering  
University of Michigan  
Ann Arbor, MI 48109

## Abstract

This extended abstract describes ongoing work in the development of an intelligent assistant that interacts with its user in a personalized fashion, deciding whether, when and how to interact based on a user's needs and preferences. I consider two types of users: people who work in an office environment and require assistance with managing their daily meeting and project schedule, and people with cognitive disabilities who require guidance in performing their daily tasks. I will examine various interaction types including reminders, requests for confirmation of a task having been completed, requests for permission to assist in the performance of a task, and requests for feedback. Both explicitly stated and implicit user preferences will be used as input to the learning mechanism, and the assistant will be evaluated based on its learning efficiency and the level of user satisfaction it achieves. This development of a preference learning model for interaction management constitutes the primary contribution of my work.

## Problem Statement

This extended abstract describes a dissertation, now in its preliminary stages, that explores preference learning within the context of an interactive, intelligent assistant. A primary challenge to an intelligent assistant is that it must adapt to meet a user's preferences while simultaneously managing that user's daily plan. In doing so, and in accordance with prior work in this area (McCarthy & Pollack 2002), it must balance a number of potentially conflicting goals. These include: (1) providing assurance that the user is aware of planned activities; (2) overseeing the effective performance of each task without introducing inefficiency; (3) maximizing the user's autonomy to reduce over-reliance; (4) performing the task or assisting the user in performing the task when appropriate; and (5) achieving a high level of user satisfaction.

Also of particular importance is the fact that the system must function in an interactive setting in which an explicit training phase is infeasible. Furthermore, a user's preferences may change over time and thus the system must perform its learning solely through its interactions with the user.

To ensure broad applicability of the results, I am targeting two very different user populations: busy executives who re-

quire assistance with meeting and project management (e.g., through a personalized assistant such as CALO (Mark & Perrault 2004)), and people with cognitive impairment who require assistance with activities of daily living (e.g., using a cognitive orthotic such as Autominder (Pollack *et al.* 2003)).

## Related Work

There are a number of interactive scheduling tools, such as CAP (Mitchell *et al.* 1994) and INCA (Gervasio, Iba, & Langley 1999), that learn a user's scheduling preferences through observation. These systems incorporate the learned preferences into a user model that can provide more accurate scheduling suggestions to the user, but they do not perform automatic calendar maintenance. Microsoft's LookOut system (Horvitz 1999) does schedule meetings autonomously but does not provide interactive prompting services.

Numerous reminding systems have been developed in the area of assistive technology for the cognitively impaired, such as Autominder (Pollack *et al.* 2003), COACH (Boger *et al.* 2005) and Neuropage (Wilson *et al.* 1997). These systems provide an automated interaction service, but they are not tailored to individual users' preferences.

There has been work in the area of adaptive user interfaces, including ARNAULD (Gajos & Weld 2005) and EIA (Liu, Wong, & Hui 2003), that explores the question of how to adapt system services to a user's preferences to create a desirable user experience. My research builds upon this prior work, extending it to consider the trade-off between preference learning and satisfactory interaction.

## Research Plan

I will begin by focusing on one particular type of interaction for an intelligent assistant: reminder generation. To this end, I am investigating the use of machine learning algorithms to support the personalization of reminders. Key research questions include: (1) How to model the world with respect to user goals, user focus of attention, agent action space; (2) How to manage the challenge of interactive systems: balancing desirable solutions with solutions that speed the rate of learning; (3) How to determine the amount of feedback that it is feasible to obtain from the user (and from different types of users); and (4) How to match the learning algorithm to the characteristics of the available feedback.

A person's preferences may be expressed explicitly or implicitly. If the system provides its user with choices, or allows for feedback or for the user to initialize the system's user model with certain preferences or goals, then that system has access to explicitly stated preferences (such as: *I prefer a voice prompt to a beeping noise*). In contrast, the system can recognize implicit user preferences by monitoring the amount of user success with respect to the interactions that the system offers its user (e.g., recognizing that the user performs a task immediately upon receiving a prompt for that task generally indicates a successful prompt).

I plan to study the manner in which a system can learn a user's explicit preferences over the type of assistance the user receives. This includes (1) allowing the user to input preferences at the onset of communication with the assistant, (2) providing the user with choices from which the user can choose the most favorable option, and (3) requesting user feedback to be incorporated into future decision making.

I have completed work in the area of active learning of scheduling preferences (Weber & Pollack 2007), in which the user of an interactive calendar program is presented with a series of scheduling options, and the user's choice of the most desirable option is used to adjust the system's user model. Similarly, in the domain of adaptive reminding, a user can be prompted in various ways for a particular task. I have also developed a simulator for testing the results of using various interaction mechanisms for learning explicit and implicit user preferences. A preliminary version of this simulation system is described in (Weber & Pollack 2005).

To test a system's learning of implicit user preferences, I will incorporate activity recognition into the learning framework and compare against a baseline interaction strategy in which reminding features (such as timing and signal type) are chosen randomly. The resulting framework will be paired with the explicit preference learning engine to create a full reminding system. Then, once I have achieved success in the reminding domain I will extend the system to handle other forms of interactions and address the issue of agent-driven adjustable autonomy: allowing the system to decide which scenarios may warrant which forms of assistance in addition to, or in lieu of, some form of passive interaction. This requires developing a model of adjustable autonomy (Bradshaw *et al.* 2004), and formalizing the problem of computing the expected risk and benefit of performing actions on behalf of the user (Berger 1985; Horvitz, Breese, & Henrion 1988).

## Research Questions

There are a number of issues that must be addressed as this research progresses. The first is the question of how to devise the proper evaluation metrics for the adaptive system, i.e., how to concretely define "success" of an activity or task such that it directly relates to the performance of the associated reminder. Will different levels of success or partial completion of a task be detectable by an activity recognition system? As listed above, it is also important to question the amount of information that we can assume can be provided by a user, as a seeding mechanism to the learning component. In other words, will a user know that he prefers an ear-

lier reminder for lunch than for his afternoon engagements? And if so, how will the system model exceptions to these preferences? Finally, it is critical to consider the situation in which a user's preferences evolve over time, and in particular how our system will account for such circumstances.

## References

- Berger, J. O. 1985. *Statistical Decision Theory and Bayesian Analysis*. New York: Springer.
- Boger, J.; Poupart, P.; Hoey, J.; Boutilier, C.; Fernie, G.; and Mihailidis, A. 2005. A decision-theoretic approach to task assistance for persons with dementia. In *Proceedings of the Int'l Joint Conference on Artificial Intelligence*.
- Bradshaw, J. M.; Feltovich, P.; Jung, H.; Kulkarni, S.; Taysom, W.; and Uszok, A. 2004. Dimensions of adjustable autonomy and mixed-initiative interaction. *Lecture Notes in Computer Science* 2969:17–39.
- Gajos, K. Z., and Weld, D. S. 2005. Preference elicitation for interface optimization. In *Proceedings of User Interface Software and Technology*.
- Gervasio, M.; Iba, W.; and Langley, P. 1999. Learning user evaluation functions for adaptive scheduling assistance. In *Proceedings of the Int'l Conference on Machine Learning*.
- Horvitz, E. J.; Breese, J. S.; and Henrion, M. 1988. Decision theory in expert systems and artificial intelligence. *Journal of Approximate Reasoning, Special Issue on Uncertainty in Artificial Intelligence* 2:247–302.
- Horvitz, E. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of Computer Human Interaction 1999*. ACM Press.
- Liu, J.; Wong, C. K.; and Hui, K. K. 2003. An adaptive user interface based on personalized learning. *Intelligent Systems* 18:52–57.
- Mark, B., and Perrault, R. C. 2004. Calo: Cognitive assistant that learns and organizes.
- McCarthy, C. E., and Pollack, M. E. 2002. A plan-based personalized cognitive orthotic. In *Sixth International Conference on AI Planning and Scheduling*.
- Mitchell, T. M.; Caruana, R.; Freitag, D.; McDermott, J.; and Zabowski, D. 1994. Experience with a learning personal assistant. *Communications of the ACM* 37:80–91.
- Pollack, M. E.; Brown, L.; Colbry, D.; McCarthy, C. E.; Orosz, C.; Peintner, B.; Ramakrishnan, S.; and Tsamardinos, I. 2003. Autominder: An intelligent cognitive orthotic system for people with memory impairment. *Robotics and Autonomous Systems* 44:273–282.
- Weber, J. S., and Pollack, M. E. 2005. Simulating users to support the design of activity management systems. In *Winter Simulation Conference*.
- Weber, J. S., and Pollack, M. E. 2007. Entropy-driven online active learning for interactive calendar management. In *Proceedings of the International Conference on Intelligent User Interfaces*.
- Wilson, B.; Evans, J.; Emslie, H.; and Malinek, V. 1997. Evaluation of neuropage: A new memory aid. *Journal of Neurology, Neurosurgery and Psychiatry* 63:113–115.