Intelligent Monitoring in a Robotic Assistant for the Elderly

Sailesh Ramakrishnan

Intelligent Systems Program
University of Pittsburgh
sai+@pitt.edu

Martha E. Pollack

Department of Computer Science and
Intelligent Systems Program
University of Pittsburgh
pollack@cs.pitt.edu

The NurseBot project is developing a mobile robot that is intended to assist elderly people suffering from mild cognitive disorders in their everyday life (see http://www.cs.cmu.edu/~nursebot). One of the main components of the project is intelligent reminding, which is useful when an elderly person has mild memory problems. The robot possesses information about the elder's daily activities, and monitors their performance, providing reminders when needed. The reminders may address activities that are critical to the elder's health and safety (taking their medications), as well as activities that support the elder's general happiness (their favorite TV program). An intelligent reminder system must be able to analyze its user's plans and compute a monitoring plan of its own.

One of the main aspects of intelligent monitoring is to identify which activities should be monitored. It is infeasible to monitor *everything*, for several reasons. First, the robot may not be able, or might choose not to follow its user around all the time, for example, it might not enter the user's bathroom. Second, the kinds of and the accuracy of the sensors may limit the ability of the robot to monitor certain activities. Third, too much monitoring may annoy the elderly user, who may be attentive to and punctual about certain activities, only requiring monitoring for other activities. Prior work on selecting activities to monitor during plan execution (e.g. Pollack and McCarthy 1999) has not focused on these aspects of the problem, but instead has distinguished between environmental changes that influence an existing plan in some way, and those that do not.

The intelligent monitoring system we are developing learns which activities need to be monitored, and when reminders need to be issued. It is initially provided with a detailed representation of the user's plans, which may be modified, via additions, deletions, or changes, as time passes. These modifications are handled by PMA, the Plan Management Agent (Pollack, Tsamardinos and Horty 1999). The representation may identify certain activities as *critical*: these must be monitored all the time (e.g. medicine-taking). Non-critical activities fall into two classes: those that should be monitored, with reminders presented as needed; and those that only need to be periodically monitored.

Our approach makes use of a Bayesian belief network that represents our belief about the likely time of the user's performance of each planned activity, contingent on his or her performance of certain earlier activities. (For instance, the time at which the user eats breakfast may depend on the time at which s/he got up.) We then run simulations on the model of the user's activities to identify which ones are likely to be late and hence should be monitored. More specifically, our algorithm involves four steps: (1) Convert the elderly person's plan from PMA's temporal network representation to a Bayesian belief network. (2) Perform numerous simulations. In each, a time interval for each activity is randomly chosen, subject to the conditional probabilities in the network. (We make use of a discretized notion of time.). (3) Use the results of the simulations to compute the probability distribution of each activity's execution over time. The probability of an activity's occurrence during time interval I is the fraction of simulations in which that interval was chosen for that activity. (4) This probability is combined with a weighting factor for that activity and is then compared to a threshold. The weighting factor for critical activities is infinitely large. An activity is flagged if the probability it will be late exceeds the threshold. Currently, we are using thresholds as a simplified mechanism for identifying activities to monitor; we plan to incorporate a more decision theoretic mechanism in a future version. Monitoring actions for the flagged activities are incorporated into the robot's plan and are executed. Each time the robot observes an activity or has information about a change, this process is repeated with the new information incorporated into the network, and if necessary the monitoring plan is modified.

References

Pollack M. E. and McCarthy, C., Towards Focused Plan Monitoring: A Technique and an Application to Mobile Robots, *IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 1999.

Pollack, M. E., Tsamardinos, I. and Horty, J. F. Adjustable Autonomy for a Plan Management Agent, *AAAI Spring Symposium on Adjustable Autonomy*, Stanford, CA, March, 1999.