

Contextual Computer Support for Human Activity

Donald J. Patterson and Dieter Fox and Henry Kautz

{djp3,fox,kautz}@cs.washington.edu

University Of Washington

Department of Computer Science and Engineering

Kenneth Fishkin and Mike Perkowitz and Matthai Philipose

{kfishkin,mikep,matthai}@intel-research.net

Intel Research

Seattle, WA USA

The modern knowledge worker has become very adept at working with desktop computers through familiar user interface devices such as keyboards, mice and screens. This interaction has relied on humans adapting to the context of the computer by going to the computer, sitting down and typing. This paradigm has been sufficient to assist people in an enormous number of tasks, but at the same time has limited the breadth of tasks in which a computer can assist.

For general-purpose computers to extend their assistance further, it is necessary for them to adapt to the human context. Laptops and personal digital assistants (PDAs) are the first wave of this adaptation in that they partially adapt to a user's location context. While miniaturization and mobilization have enabled this adaptation and provided new functionality, these devices generally remain blind to the world around them. As a result they have become intrusive in our daily lives as cell-phones ring in movie theaters, laptops interrupt presentations with instant messages, and PDAs rattle and beep for our attention in meetings. However, the goal of having computers fully aware of their user's context is difficult and broad.

Intel Research Seattle (IRS) (Intel Research Seattle 2003) and the Assisted Cognition project (Kautz *et al.* 2002b; 2002a) at the University of Washington in Seattle are developing a system entitled Guide to learn a small component of context by probabilistically inferring human activity. There are three research hypotheses that drive Guide's development:

- A computer can detect the performance of high-level, complex activities by directly sensing user's manipulation of objects involved in those activities
- Rich and comprehensive libraries of activities can be created by datamining unstructured text.
- Modern stochastic inference algorithms can effectively handle complex, high-dimensional models of human behavior.

Guide is designed to be scalable to many thousands of activities, to be robust, and to be adaptive to new behaviors, and would be valuable simply for its ability to record diaries of daily activities. This would, for example, greatly assist the monitoring of Activities of Daily Living (ADLs) (Lawton 1990; 1983); a task that currently burdens medical caregivers of the elderly. Eventually such a system would be useful as a basis for providing information with more contextual sensitivity, for giving proactive advice to avoid safety and security errors and for training workers in the best known methods of manufacturing operations.

Our implementation relies heavily on Radio Frequency Identification (RFID) tags. RFID tags are small computer chips embedded in a sticker the size of a postage stamp. There is no power supply associated with an RFID tag, but when a special antenna broadcasts an information request, the tag collects the energy from the request and re-broadcasts a uniquely identifying number that is embedded within it. The querying antenna listens for the id numbers of any tags which respond and processes the information. Versions of this technology are familiar from their use in retail anti-shopping-lifting systems. RFID tags are cheap, durable, low-maintenance and have very little noise associated with their localization.

We have prototyped an antenna embedded in a glove which is able to see all tags that are located within three inches of the palm of the wearer. This closely approximates sensing when an object is touched and is compelling because it does not require external infrastructure, has a small form factor and effectively guards user privacy with a single point of information control. Guide uses information obtained from this RFID sensor stream, coupled with models of human activity and a powerful inference engine, to probabilistically reason about likely activities in the environment. Once a sensor stream has been recorded, the system is then able to update its models using unsupervised machine learning techniques developed in the context of transportation activity inference (Patterson *et al.* 2003).

The system relies on three software components to ac-

compish this, a data miner, an activity inference engine and a visualizer. The data mining component is critical to achieving scalability of the system. Given the wide range of human activity it would be problematic to require an expert to generate a model of every interesting activity that a person could perform. Instead we model ordered compound activities by mining “how-to” sites (e.g. (All Recipes 2003; eHow Inc. 2003)) and interpret the textual order of the steps as sequential constraints. We augment these sequential models with prior probabilities of object touches, given an activity, by leveraging features of the Google programming API (Google Labs 2003). To date, we have mined the sequential structure of roughly fifteen thousand activities.

The activity inference engine accepts the models from the data miner, converts them to Dynamic Bayesian Networks (Murphy 2002) and uses particle filters (Fox *et al.* 2001; Doucet, de Freitas, & Gordon 2001) to infer activities given an RFID sensor stream. The visualizer gives insight into the reasoning process by showing how the sensors affect the inference engine’s belief state and allow a user to improve and debug the generated models. By using Expectation-Maximization the activity inference engine also has the capability to learn model parameters in an unsupervised manner.

As a preliminary experiment we asked 14 individuals to wear the RFID glove and perform ADLS in a home instrumented with RFID tags. We asked each subject to accomplish 12 out of 14 goals in the way that they normally would and record the sequence. Example goals included making a snack, making a hot beverage, and taking care of personal appearance. Under these circumstances and using hand-made models, the system performs well, correctly recognizing goals with 88% precision and 73% recall. (Philipose *et al.* 2003)

As we continue to integrate and extend the components of our system, we are actively seeking input about how to best support the proactive software components that we regard as the consumers of our reasoning. Understanding the applications and needs of these components will further drive our research agenda and inform our development efforts.

References

- All Recipes. 2003. All Recipes website. <http://www.allrecipes.com/>.
- Doucet, A.; de Freitas, N.; and Gordon, N., eds. 2001. *Sequential Monte Carlo in Practice*. New York: Springer-Verlag.
- eHow Inc. 2003. eHow website. <http://www.ehow.com/>.
- Fox, D.; Thrun, S.; Dellaert, F.; and Burgard, W. 2001. Particle filters for mobile robot localization. In Doucet *et al.* (2001).
- Google Labs. 2003. Google Web APIs. <http://www.google.com/apis/>.
- Intel Research Seattle. 2003. Intel Research Seattle Guide Project Website. <http://seattleweb.intel-research.net/projects/guide/>.

Kautz, H.; Arnstein, L.; Borriello, G.; Etzioni, O.; and Fox, D. 2002a. The Assisted Cognition Project. In *Proceedings of UbiCog '02: First International Workshop on Ubiquitous Computing for Cognitive Aids*.

Kautz, H.; Fox, D.; Etzioni, O.; Borriello, G.; and Arnstein, L. 2002b. An Overview of the Assisted Cognition Project. <http://www.cs.washington.edu/homes/kautz/ac/acaaai02.pdf>.

Lawton, M. P. 1983. Environment and Other Determinants of Well-Being in Older People. *The Gerontologist* 23(4):349–357. NLM:0375327;PMID:6352420.

Lawton, M. P. 1990. Aging and Performance of Home Tasks. *Human Factors* 32(5):527–536. NLM:0374660;PMID:2074107.

Murphy, K. 2002. *Dynamic Bayesian Networks: Representation, Inference and Learning*. Ph.D. Dissertation, UC Berkeley, Computer Science Division.

Patterson, D. J.; Liao, L.; Fox, D.; and Kautz, H. 2003. Inferring High-Level Behavior from Low-Level Sensors. In *Proceedings of UBIComp 2003: The Fifth International Conference on Ubiquitous Computing*.

Philipose, M.; Fishkin, K.; Perkowitz, M.; Patterson, D.; and Hähnel, D. 2003. The Probabilistic Activity Toolkit: Towards Enabling Activity-Aware Computer Interfaces. Technical report, Intel Research Seattle. http://www.cs.washington.edu/homes/djp3/AI/AssistedCognition/publications/Intel/irs_tr_03_013.doc.