Guide: Towards Understanding Daily Life via Auto-Identification and Statistical Analysis

Kenneth P. Fishkin[1], Henry Kautz[2], Donald Patterson[2], Mike Perkowitz[2], Matthai Philipose[1]

[1] Intel Research Seattle, 1100 NE 45th St, 6th floor
Seattle, WA 98105, USA
{kenneth.p.fishkin, matthai.philipose}@intel.com

[2] Department of Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350, USA
{kautz, djp3, mikep}@cs.washington.edu

http://seattleweb.intel-research.net/projects/guide

Abstract. Many recent studies have underscored the applicability to healthcare of a system able to observe and understand day-to-day human activities. The Guide project is aimed at building just such a system. The project combines novel sensing technology, expressive but scalable learners and unsupervised mining of activity models from the web to address the problem. An early prototype, Carnac, has shown considerable promise. This paper provides a high-level overview of the Guide approach, briefly discusses Carnac, and details our expectations of the workshop.

1 Introduction

Understanding the day-to-day activities of individuals (such as eating, cleaning, cooking, watching television, exercising, shaving, and playing) has been recognized as a capability with a wide range of applications to healthcare. Applications include reminders to perform missed activities, prompts to help in completing activities, monitoring to assess the ability to live independently, notifications to caregivers as a response to anomalous or undesirable patterns of activity and logs to help understand a patient's current state based on recent activity [2, 5, 6, 7, 9, 10, 12]. Many systems address the problem of activity recognition, but all the ones we are aware of are severely limited in the variety of activities they recognize, the robustness and speed with which they recognize them, and/or the ease of adding new activities for the system to detect. The Guide project is a general approach to the activity recognition problem that is intended to address these limitations. The goal is to build a system (sensors and learners) that can be practically deployed and used by consumers to detect activities such as the ones listed above.

2 Current work

We sketch below how Guide-based activity detectors may be used. We then outline the key interesting components of Guide.

2.1 Usage model

Guide assumes that all “interesting ” objects are tagged with postage-stamp sized 10-cent transceivers called Radio Frequency Identification (RFID) [4] tags. These tags contain circuitry that harvest power from interrogating readers, and use this power to send back a unique identifier to the reader. In our current deployments, we have tagged objects such as utensils, furniture, clothes, tools and appliance controls in an actual home without affecting the normal usage of these objects. A person who wants their activities to be tracked wears a wrist-worn or glove RFID reader. As the person then proceeds about their daily activities and touches tagged objects (actually, places their hand within a few inches of the tags), the system observes the sequence of objects touched and deduces the nature of the activity currently taking place.

A client application of the system can specify activities of interest as English strings (such as “playing the violin” or “doing the dishes”) to the system. Over time, the system provides the client with likelihoods for the specified activities.

2.2 Key Components

Our approach has four main components of interest.

First, we use a new class of sensors, that we call *direct association* sensors, which can directly attribute sensed properties to individual objects. Concretely, these sensors are variations of RFID tags. Direct association allows us to monitor tens of thousands of objects in a home-sized space for properties such as touching, proximity and movement. We have developed wrist-worn and glove-like RFID-based sensors capable of sensing which tagged objects have been touched, and a mobile-robot mounted reader capable of localizing tags in an unstructured space to a resolution of less than a meter. Work to detect moving tags from signal strength variation is in progress.

Second, we model activities coarsely in terms of the objects involved in each activity. Involvement may include being touched, being moved, or being close to the person performing the activity, and can be sensed directly by our sensors. For example, the activity of “shaving” may with high probability involve touching objects “razor” and “shaving cream”. We model the dependence between activities and object-involvement with Bayesian belief net representations, using dynamic formulations [8] as appropriate.

Third, we seek to select from thousands and tens of thousands of activity models at any given time, and eventually to track of the order of ten simultaneous activities, all preferably in real time. Our basic technique for achieving scalability is to rely on the fact that we have large quantities of reliable sensed data about activities: we keep the structure of our belief networks extremely simple while using as large a variety of observations as possible to identify activities. For instance, our initial prototype used Bayes Nets with a simple bipartite structure: they mapped each activity to the objects involved with that activity (along with the probability of involvement). Although even this very simple model is surprisingly effective, it has become clear that capturing ordering constraints among sub-activities and timing constraints is useful. In a closely related project [11], we are now exploring timed, dynamic, relational variants of Bayes Nets as the representation of choice.

Fourth, we employ data mining techniques to extract prior models for activities directly from the web in an unsupervised fashion. Our current techniques allow us to automatically mine simple bipartite Bayes Net models (such as those describe above) for *arbitrary* activities if we do not seek to model temporal structure. One surprisingly good technique for mining the prior probability of an activity A given objects $O_{1...}O_n$ is the following. Let n_1 and n_2 be the number of pages in the web that match strings "A" and "$A\ O_{1...}O_n$" respectively (say via Google [5]). Then $Pr(O_{1...}O_n|A) = n_2/\ n_1$. When we are interested in modeling ordered sub-activities of compound activities, we mine "how-to" sites [1, 3]. We interpret the textual order of the how-to steps as possible temporal constraints. We have mined the temporal structure of roughly fifteen thousand activities (including *e.g.* those for boiling an egg, and cleaning a bathtub).

Combining some of these components has resulted in an early prototype activity recognition system named Carnac. Carnac attempts to understand routine human activity being performed in unstructured spaces. It is able to assign likelihoods to arbitrary activities at low latency. We provide Carnac with a list of roughly fifty activities to choose between (including many standard activities of daily living such as brushing teeth and eating cereal), and based on observations Carnac assigns a likelihood score to the activities. Carnac's perception of likely ongoing activities has been surprisingly in line with our intuition.

3 Expectations from workshop

We would like to understand better how both the components (*e.g.* sensors) and the end products of a Guide-style system can fit into the agendas of stakeholders in computer-aided healthcare. In particular, we would like to integrate Guide-style activity detectors into practical end-user applications. We are therefore interested in being learning about, and possibly collaborating with, researchers who have a need for an activity detector as part of their system.

We have selected the problem of automatically filling in Activities of Daily Living (ADL) forms as a test case for Guide. Although we have an expensive list of these activities, we would like to talk to field researchers who have an idea of the pragmatics surrounding the collection and use of ADLs.

We see the RFID-based sensing as complementary to other sensing modalities. For instance, it is plausible that we can use the context provided by Guide to significantly improve the performance of computer vision based systems. Such systems, in turn, would be able to reason about aspects of the activity (for instance, taking into account constraints on geometry and color) that Guide cannot. We would therefore like to communicate with researchers applying such complementary technologies to the problem of activity understanding.

Finally, we are interested in deploying Guide. We would like to talk to field researchers who have the opportunity and the desire to deploy such technology.

Biographies

Ken Fishkin is a researcher at Intel Research Seattle, and an Affiliate Professor of Computer Science at the University of Washington. He has over 15 years experience in Industrial R&D, including stints at Pixar (4 years), and Xerox PARC (8 years). While at PARC, he co-founded the "eXtreme UI" research effort, which helped invent what are now called "Tangible User Interfaces". He has published widely in the fields of Computer Graphics and User Interfaces. His current research interests center on Ubiquituous Computing ("Ubicomp") user interfaces and systems. He holds Bachelor's degrees in Computer Science and in Mathematics from the University of Wisconsin (Madison) , and a Master's Degree in Computer Science from the University of California (Berkeley).

Henry Kautz is an Associate Professor in the Department of Computer Science and Engineering at the University of Washington. He joined the faculty in the summer of the year 2000 after a career at Bell Labs and AT&T Laboratories, where he was Head of the AI Principles Research Department. His academic degrees include an A.B. in mathematics from Cornell University, an M.A. in Creative Writing from the Johns Hopkins University, an M.Sc. in Computer Science from the University of Toronto, and a Ph.D. in computer science from the University of Rochester. He is a recipient of the Computers and Thought Award from the International Joint Conference on Artificial Intelligence and a Fellow of the American Association for Artificial Intelligence. In 1998 he was elected to the Executive Council of AAAI, and in 2000 was Program Chair for the AAAI National Conference.

Donald J. Patterson is a Ph.D. candidate in the Computer Science Department at the University of Washington. He graduated from Cornell University with a B.S. in Computer Science and an M.Eng. in Electrical Engineering. After a spending a few years in the military, he returned to academia as an NDSEG fellow and is currently researching artificial intelligence and its application to health care through the Assisted Cognition Project. He is spending the summer working on ubiquitous computing applications at Intel Research Seattle.

Mike Perkowitz is currently working as a postdoctoral researcher at the University of Washington. His current interest is in applying machine learning techniques for to ubiquitous computing problems. He spent two years as senior engineer at Appliant, a web-performance measurement company. His academic degrees include a Ph. D. in Computer Science from the University of Washington and a B.S. in Cognitive Science from Brown University.

Matthai Philipose is a researcher at Intel Research Seattle. His chief interest is in making ubiquitous computing systems useful and practical. Currently, he co-leads the Guide project (which is devoted to acquiring and understanding auto-identification data), and leads the River project (which is devoted to understanding how to make ubiquitous computing systems easier to manage). He did his graduate work at the University of Washington, and his undergraduate work at Cornell University.

References

1. All Recipes website, http://www.allrecipes.com
2. Rich DeVaul, Steve Schwatz, and Sandy Pentland. The memory glasses project. online at http://www.media.mit.edu/wearables/mithril/memory-glasses.html, 2000.
3. eHow website, http://www.ehow.com
4. Finkenzeller, K. *RFID Handbook*, John Wiley and Son, Ltd., New York 1998.
5. Google website, http://www.google.com
6. Intel Proactive Health group. Proactive Health Case Study: Aging in Place. Online at http://www.intel.com/research/prohealth/cs-aging_in_place.htm
7. Henry Kautz, Oren Etzioni, Dieter Fox, Dan Weld, and Lokendra Shastri. Foundations of Assisted Cognition. University of Washington Computer Science and Engineering Technical Report CSE-03-AC-01, March 2003.
8. K. Murphy. Dynamic Bayesian Networks: Representation, Inference and Learning. PhD thesis, UC Berkeley, Computer Science Division, 2002.
9. C.E. McCarthy and M. Pollack. A plan-based personalized cognitive orthotic. In *Proceedings of the 6th International Conference on AI Planning and Scheduling (AIPS 2002*), 2002.
10. E. Mynatt, I. Essa, and W. Rogers. Increasing the opportunities for aging-in-place. In *Proceedings of ACM Conference on Universal Usability*, 2000.
11. D. Patterson, D. Fox, H. Kautz, M. Philipose. Learning Activities of Daily Living Using RFID Tags. Submitted to this workshop.
12. M. E. Pollack, C. E. McCarthy, S. Ramakrishnan, I. Tsamardinos, L. Brown, S. Carrion, D. Colbry, C. Orosz, and B. Peintner. Autominder: A planning, monitoring, and reminding assistive agent. In *Proceedings of the 7th International Conference on Intelligent Autonomous Systems*, 2002.