

Inferring Activities from Interactions with Objects

Recognizing and recording activities of daily living is a significant problem in elder care. A new paradigm for ADL inferencing leverages radio-frequency-identification technology, data mining, and a probabilistic inference engine to recognize ADLs, based on the objects people use.

A key aspect of pervasive computing is using computers and sensor networks to effectively and unobtrusively infer users' behavior in their environment. This includes inferring which activity users are performing, how they're performing it, and its current stage. The elder-

care field is a prime, yet difficult application area for inferring whether and how people with early-stage cognitive decline are performing *activities of daily living*.¹ (For more on ADLs, see the "Activities of Daily Living" sidebar.)

Recognizing ADLs, particularly in the home, is challenging on several fronts. First, because users can perform ADLs in various ways, models of activities

and recognition software must adapt to this variety. Second, the underlying sensors must report the features required of them robustly across various sensing contexts (such as light levels, sound levels, and locations). Third, given the large number of ADLs—20 to 30 classes (such as making a meal) with thousands of instances—a system should model each activity with minimum human effort. Addressing these challenges simultaneously has been a key barrier to success for ADL-monitoring systems.

We propose an approach that addresses these

challenges and shows promise in automating some types of ADL monitoring. Our key observation is that the sequence of objects a person uses while performing an ADL robustly characterizes both the ADL's identity and the quality of its execution. So, our Proactive Activity Toolkit (PROACT)

- Represents activities as a probabilistic sequence of objects used
- Adapts a cheap, durable, easy-to-use sensing technology to robustly sense the objects being used across various sensing and use contexts
- Mines probabilistic models of activity use from plain English descriptions of activities, such as recipes

Proact

Our system has three components: specialized sensors to detect object interactions, a probabilistic engine that infers activities given observations from sensors, and a model creator that lets us easily create probabilistic models of activities. (The "Related Work in ADL Inferencing" sidebar describes related research.)

Sensors

We tag objects of interest using radio-frequency-identification (RFID) tags, which we can attach unobtrusively onto objects as small as a plastic fork. These tags are postage stamp sized, durable, battery free, and inexpensive (US\$0.40 each and falling). When interrogated by a reader,

Matthai Philipose, Kenneth P. Fishkin, and Mike Perkowitz
Intel Research Seattle

Donald J. Patterson, Dieter Fox, and Henry Kautz
University of Washington

Dirk Hähnel
University of Freiburg

Activities of Daily Living

The health-care community distinguishes between *activities of daily living* (ADLs), which don't involve interactions with instruments (such as telephones and appliances),² and *instrumental activities of daily living* (IADLs), which do.¹ However, the standard practice (which we follow) is to use ADL generally to refer to both sets.

ADL monitoring has become a common, valued technological goal for three reasons. First, ADL monitoring is an important activity in health care. For example, in the US, any nursing home that receives Medicare funds must record and report ADLs. Trained caregivers spend considerable time measuring and tracking ADL accomplishment. This monitoring is time consuming, error prone, and invasive.³ Automated aids that can augment caregiver work practice are of great interest. We stress that our inferencing isn't designed to replace caregivers but to let them focus their attention on caregiving.

Second, ADLs are general—they're common activities that people perform daily and have potential interest outside the elder-care field.

Finally, home ADLs are challenging to recognize owing to the

number of activities people perform, the number of ways they perform those activities, and the difficulty of deploying sensing technology (as opposed to, for example, measuring office computer activity). ADL recognition also has a higher burden for maintaining privacy owing to the intimate nature of the activities and environment.

REFERENCES

1. S. Katz et al., "Studies of Illness in the Aged. The Index of ADL: A Standardized Measure of Biological and Psychological Function," *J. Am. Medical Assoc.*, vol. 185, 1963, pp. 914–919.
2. S. Katz, "Assessing Self-Maintenance: Activities of Daily Living, Mobility, and Instrumental Activities of Daily Living," *J. Am. Geriatrics Soc.*, vol. 31, no. 12, 1983, pp. 721–726.
3. A. McDonald and J. Curtis, *Not Alone: A Good Practice Guide to Working with People with Dementia in Sheltered Housing*, Anchor Trust, 2001.

Figure 1. An RFID-detecting glove.

they respond with a globally unique identifier. Researchers have investigated the limited use of RFID-tagged objects in user interfaces before,² but they're now so cheap and the readers are so small that qualitatively different use scenarios are feasible. RFID tags' durability and their ability to be placed out of sight make them particularly well suited for ADL monitoring of elders with early-stage dementia, who tend to disassemble foreign objects.

The tags' greatest advantage is that they naturally report data at the object level. The best long-term solution will be an integrated approach that augments RFID information with other sensor streams to fill in the gaps—for example, an audio monitor for analyzing a baby's noises. Here, we report on an RFID-only approach, with the goal of understanding its features and shortcomings.

Even with arbitrarily cheap tags, no one is going to tag every sock or ice cube. We can achieve the needed functionality by tagging dispensers and containers such as a sock drawer or ice tray. The



more tagged objects, the more powerful the inferencing and the more detailed the caregiver's report. The fewer tagged objects, the lower the cost and the shorter the deployment time. Any particular house or user can assess that trade-off as desired. In addition, we can boost a system later by adding more tags at points of interest if and as desired. Tagging an object involves sticking an

RFID tag on it and making a database entry mapping the tag ID to a name. Emerging standards might soon make the database entry unnecessary (see the EPC Global Web site, www.epcglobalinc.org).

To sense tags, we used a glove-based RFID reader, built from commercial parts. The user wears a prototype glove with an RFID-detecting antenna in the palm (see Figure 1). The antenna is

Related Work in ADL Inferencing

A number of researchers have investigated ADL inferencing. Alex Mihailidis, Geoff Fernie, and Joseph Barbenel successfully used cameras and a bracelet to infer hand washing.¹ Dadong Wan used radio-frequency-identification (RFID) tags functionally as contact switches to infer when users took medication.² Tracy Barger and her colleagues used contact switches, temperature switches, and pressure sensors to infer meal preparation.³ Quan Tran, Khai Truong, and Elizabeth Mynatt used cameras to infer meal preparation.⁴ Anthony Glascock and David Kutzik used motion and contact sensors, combined with a custom-built medication pad, to get rough inference on meal preparation, toileting, taking medication, and up-and-around transference.⁵ Ilkka Korhonen, Paula Paavilainen, and Antti Särelä used a custom wearable computer with accelerometers, temperature sensors, and conductivity sensors to infer activity level.⁶ Michael Mozer used 13 sensors to infer home energy use, focusing on the heating-use activity.⁷ Eric Campo and Marie Chan used motion detectors to infer rough location.⁸ Valerie Guralnik and Karen Haigh used motion sensors, pressure pads, door latch sensors, and toilet flush sensors to infer behavior.⁹

These systems have all had to perform high-level inferencing from low-level, coarse sensor data reporting. Some have added special pieces of hardware to help surmount this, but progress toward rigorous ADL detection has nevertheless been slow. Only a few researchers have reported the results of any preliminary user testing,^{1,5,8,9} and all but one report¹ are anecdotal. The level of inferencing has often been limited—for example, reporting only that a person entered the living room and spent time there. Although each of these difficult investigations has been valuable, they lack interconnection and generality. Research tuned to detecting hand washing or tooth brushing have had nearly no synergy, each using its own set of idiosyncratic sensors and algorithms on those sensors. A home deployment kit designed to support all these ADLs would be a mass of incompatible and noncommunicative widgets. Our approach instead focuses on a general inferencing and sensor capability that we can use for many ADLs. Although initial performance will be worse than tailored algorithms for a particular activity, our approach applies across multiple ADLs, and we can tune it with experience and use additional sensors as needed to improve its performance.

A new approach, similar in spirit to ours, is the research on MIT's House_n project (http://architecture.mit.edu/house_n). Like us, this project places a single type of object-based adhesive sensor in

structurally unmodified homes. Sensor readings are later analyzed for various applications—kitchen design, context sampling, and potentially ADL monitoring. These researchers share our contention that fine-grained measurement of object use is a good indicator of activity. Our work differs from theirs in two key ways. First, we adapt sensors that we can attach unobtrusively to objects as small as spoons and toothbrushes and that require no battery; they use battery-powered sensors more suited for tagging larger objects such as cupboards and stoves. Second, we let programmers describe activities in ordinary English and provide automatic means to convert these descriptions into formal models.

REFERENCES

1. A. Mihailidis, G. Fernie, and J.C. Barbenel, "The Use of Artificial Intelligence in the Design of an Intelligent Cognitive Orthosis for People with Dementia," *Assistive Technology*, vol. 13, no. 1, 2001, pp. 23–39.
2. D. Wan, "Magic Medicine Cabinet: A Situated Portal for Consumer Healthcare," *Proc. 1st Int'l Symp. Handheld and Ubiquitous Computing (HUC 99)*, LNCS 1707, Springer-Verlag, pp. 352–355.
3. T. Barger et al., "Objective Remote Assessment of Activities of Daily Living: Analysis of Meal Preparation Patterns," poster presentation, Medical Automation Research Center, Univ. of Virginia Health System, 2002.
4. Q. Tran, K. Truong, and E. Mynatt, "Cook's Collage: Recovering from Interruptions," demo at 3rd Int'l Conf. Ubiquitous Computing (UbiComp 2001), 2001.
5. A. Glascock and D. Kutzik, "Behavioral Telemedicine: A New Approach to the Continuous Nonintrusive Monitoring of Activities of Daily Living," *Telemedicine J.*, vol. 6, no. 1, 2000, pp. 33–44.
6. I. Korhonen, P. Paavilainen, and A. Särelä, "Application of Ubiquitous Computing Technologies for Support of Independent Living of the Elderly in Real Life Settings," *Proc. UbiHealth 2003: 2nd Int'l Workshop Ubiquitous Computing for Pervasive Healthcare Applications*, 2003, www.healthcare.pervasive.dk/ubicomp2003/papers/Final_Papers/2.pdf.
7. M. Mozer, "The Neural Network House: An Environment That Adapts to Its Inhabitants," *Proc. AAAI Spring Symp. Intelligent Environments*, tech. report SS-98-02, AAAI Press, 1998, pp. 110–114.
8. E. Campo and M. Chan, "Detecting Abnormal Behavior by Real-Time Monitoring of Patients," *Proc. AAAI Workshop Automation as Caregiver*, AAAI Press, 2002, pp. 8–12.
9. V. Guralnik and K. Haigh, "Learning Models of Human Behaviour with Sequential Patterns," *Proc. AAAI Workshop Automation as Caregiver*, AAAI Press, 2002, pp. 24–30.

connected to a SkyeTek RFID reader, packaged with a Crossbow Mica Mote radio, a USB-based power supply, and a rechargeable battery. All components

except the antenna are housed in the small box on the glove. The reader samples the environment twice a second; any RFID seen is broadcast to an HP iPaq

5400, used as a wearable computer. The iPaq either stores the data onboard for future analysis or forwards it via Wi-Fi to the inference engine running on a

workstation. The reader lasts for two hours at this duty cycle.

Although for many applications and users, using such a glove would be problematic, we believe this problem is temporary. Three years ago, Albrecht Schmidt, Hans-Werner Gellersen, and Christian Merz proposed a portable RFID glove as a user interface that measured roughly 160 cm³ and wasn't wireless.² Our latest version is a little more than 30 cm³ and wireless. In a few years, portable RFID readers will likely fit into a large wristwatch or bracelet and will be feasible for ADL monitoring.

Activity models: Structure and inference

PROACT represents activities as linear sequences of activity stages. For example, we can model the making of tea as a three-stage activity: boil the water; steep the tea in the water; and flavor the tea with milk, sugar, or lemon. We annotate each stage with the objects involved and the probability of their involvement. For example, in the first stage, we might expect to see a kettle and water, with high probability. In the third stage, we might expect to see sugar, lemon, or milk, but each with a lower probability. The probabilities need not sum to one. This probability combines three sources of ambiguity: sensor error (for example, the tag is missed), model error (for example, some objects aren't known to the model), and modeling generality (for example, an object is sufficient but not necessary). Each stage might also optionally have a time to completion modeled as a Gaussian probability distribution. For example, we might indicate that the boiling stage takes on average three minutes with a standard deviation of two minutes.

The probabilistic engine converts these activity models into *dynamic Bayesian networks*. By regarding the current sub-activity as a hidden variable, and the set of objects seen and time elapsed as

observed variables, we can probabilistically estimate the activities from sensor data with the general technique of Bayes filtering. We use a type of sequential Monte Carlo approximation to probabilistically solve for the most likely activities. To solve the activity models we've sketched, we could use simpler hidden-Markov-model-based inference. However, using the more complex DBN infrastructure lets us explore complications such as partial (as opposed to total) order between activity steps and a somewhat nonstandard (but natural) interpretation of time to completion. (Further details appear elsewhere.³)

Model creation

Activity inferencing requires a model

Once we specify the objects involved in each stage, we must define the object *involvement probabilities*. Intuitively, these describe the probability of using the object in that activity state.

that defines an activity: how many stages it has and how sensor data determines the transitions between stages. These models are typically difficult to create without intensive hand tuning, owing again to the low-level nature of the input that guides the state transitions—for example, a change in the pattern of accelerometer input.

Because our raw sensor data is high level (the names of objects), we can transform high-level activity descriptions into formal models. In particular, the programmer provides a step-by-step English description of the activity in the format commonly used for specifying recipes or how-tos. We use lightweight natural-language-processing techniques to convert each step into a stage in our activity model, and we use the set of objects mentioned in the step as the set

of objects involved in that activity. We've also mined existing recipes and how-tos using this technique; details appear elsewhere.⁴

Once we specify the objects involved in each stage, we must define the object *involvement probabilities*. Intuitively, these describe the probability of using the object in that activity state. For example, making tea always involves tea and water but only occasionally involves sugar. Requiring humans to specify these probabilities is time consuming (consider all the potential objects involved in making a meal) and difficult (what percentage of sandwiches involve ham?).

Instead, PROACT automatically determines these probabilities. Our method relies on a *mirror assumption*: if an activ-

ity name co-occurs often with some object name in human discourse, then the activity will likely involve the object in the physical world. Our approach is in the spirit of previous research that used word associations on the Web to resolve ambiguities and determine world relevancies.⁵ We use them as a model of relatedness in human activity. We postulated that if an activity A occurs on n_1 Web pages (the best extant approximation of human discourse), and n_2 pages contain the activity and an object T , then the involvement probability for T in A is approximately equal to n_2/n_1 . We obtain these numbers via the Google API. This assumption can fail, but in our experience it suffices to generate a good set of probabilities, especially as we need only the terms' relative, not absolute, frequencies. We used the probabilities

Table 1
The tested activities of daily living.

ADL	Task sheet description	Notes	
1	Personal appearance	Please touch up your personal appearance.	
2	Oral hygiene	Please take care of your oral hygiene as if you were about to go to bed.	
3	Toileting	Please use the toilet.	
4	Washing	Please wash up.	
5	Housework	Please clean something.	
6	Safe use of appliances	Please use an appliance.	
7	Heating use	Please adjust the thermostat.	The thermostat was the only heating control in the house
8	Caring for clothes and linen	Please use the washer and/or dryer.	
9	Preparing a simple snack	Please prepare something to eat.	
10	Preparing a simple drink	Please make yourself a drink.	
11	Telephone use	Please use the telephone to get your horoscope.	We focused the task to encourage a short call.
12	Leisure activity	Please watch TV, read, play cards, or listen to music for a few minutes.	
13	Infant care	Please care for the baby.	We used a life-size, untagged doll.
14	Taking medication	Please take some pills.	We used candy instead of pills.

unchanged in our experiment. (A more detailed exploration appears elsewhere.⁴) The probabilities can evolve with time in response to learning or can be tuned to a particular elder (that is, for some, the percentage of sandwiches with ham is 0).

This framework, although naturally suited for object interaction reports, isn't limited to them. For example, it can easily accommodate location reports by using location words and specifying their involvement probabilities. This lets the engine realize that, for instance, touching toilet paper in the kitchen more likely indicates unpacking than toileting.

Testing the concept

The prototype nature of the RFID detector makes a true deployment with elders in cognitive decline presently infeasible. However, we can test the approach's validity by asking a slightly different question: If we could reliably detect interactions with RFID-tagged objects, could we infer ADLs? We did this test. Although PROACT can accommodate other types of sensor data, for this test we restricted ourselves to object interaction reports.

Choosing the ADLs

Of the more than 20 ADLs that caregivers can monitor, we chose 14 for evaluation (see Table 1). We eliminated the rest because the subjects in our experiment performed their tasks in another person's home; therefore, we excluded ADLs such as bathing or shopping. The 14 ADLs we tested are, to our knowledge, 11 more than any other system has attempted.

We took this set of activities, with all their vagueness and overlap (consider personal appearance versus oral hygiene versus washing up, for example) as a set of givens. We didn't narrow or constrain the activity definitions.

Procedure

We wanted to use as realistic a home setting as possible. In a perfect world, we would have instrumented the houses of each of our subjects, but this was impractical. Instead, we chose a house inhabited by one of the authors, his wife, and their two-year-old child. Using a single house has disadvantages. Subjects weren't in their own homes, and we couldn't do a longitudinal study. However, we feel this was a valid compro-

mise. Several subjects mentioned that because they weren't observed and were in a real house, they felt relaxed and at ease, and therefore acted more naturally. We instrumented the home with 108 tags in a few hours. We did this before the activity models existed to minimize biasing our tagging with key objects. We tagged as many objects as possible so that the subjects wouldn't feel steered to a narrow set.

Over the next six weeks, we tested the concept by having 14 subjects (3 male, 11 female) perform ADLs, wearing the glove in Figure 2. Each spent roughly 45 minutes in the house. Ages ranged from 25 to 63, with a mean of 39. We recruited subjects by word of mouth and paid them \$20 on completion.

After we took the subjects on a tour of the house, described the experiment, and demonstrated the glove, we gave the subjects 14 task sheets, one for each ADL. The task sheets had pictures of where the subjects could find some objects in the house to avoid unnecessary searching. We kept activity descriptions (see Table 1) broad and as close to the medical description as possible.

Subjects randomly selected 12 of the

Table 2
Experiment results

ADL no.	True positives	False positives	False negatives	Precision (%)	Recall (%)
1	11	1	1	92	92
2	7	3	2	70	78
3	8	3	3	73	73
4	3	0	6	100	33
5	8	0	4	100	75
6*	21	4	6	84	78
7	8	0	3	100	73
8	7	0	2	100	78
9	6	2	4	75	60
10	9	5	5	64	64
11	11	0	3	100	79
12	7	0	5	100	58
13	13	0	1	100	93
14	9	0	2	100	82
Total	128	18	47	88	73

*The “safe use of an appliance” task is also an optional subtask in many other tasks—for example, using the microwave to boil water for a drink or using the clothes dryer. We count the use of an appliance during another task as a positive for the “safe use of an appliance” task.

14 tasks. We then asked them if they wished to withdraw or felt uncomfortable with what they were asked to do—none did. They then went into the house and performed their 12 tasks, in any order, without observation. We waited outside. As they touched tags, they heard a beep when PROACT detected the touch. Because this round of experiments wasn’t meant to test the glove’s efficacy, we asked the subjects to retouch something if they didn’t hear a beep. This was sometimes necessary owing to the glove’s prototypical nature; typically, subjects touched approximately 60 tags, one of which they needed to repeat. While performing tasks, the subjects wrote on a sheet of paper which task they were doing. After leaving the house, the subjects gave one experimenter the sheet. We kept this sheet separate from the other experimenters until after we processed the results.

The tags remained in place for six weeks in the house, which was permanently occupied and used by an extended family (including visiting relatives). Except for a few tags that we had awkwardly placed on edges, all tags stayed usable throughout, surviving the full destructive power of a two-year-old in his native habitat.

Results

We gave the object-touch sequence for each subject (without manually segmenting between activities), along with models for the 14 activities, to the inference engine. The engine returned the most likely activity sequence of that would explain the readings. Running the activity inference took much less time than performing the activities. So, although we used PROACT offline, it can infer these ADLs in real time.

Although for this experiment we used the inference engine to report the most probable ADL at any given time, the engine has much more detailed infor-

mation. Using a probabilistic system such as this, and with object-based input, we can provide a richer set of data to a caregiver than a simple binary yes or no for a particular ADL. We can report our confidence in that conclusion and the evidence behind it. The caregiver can use this as part of his or her assessment—for example, noting that an elder is wearing the same clothes each day, that the elder is taking longer and longer to make lunch, and so forth. The more tags in the house, the richer this information.

We compared the PROACT report with the activity sequence that the subjects reported immediately after the experiment. We treat the latter as ground truth. When PROACT correctly claimed an activity occurred, it scored a true positive (TP); an incorrect claim scored a false positive (FP). If an activity occurred and PROACT didn’t report it, PROACT scored a false negative (FN). Table 2 shows the results for each ADL. We then used two standard metrics to summarize PROACT’s effectiveness. *Precision* is the probability that a given inference about that activity is correct: $TP/(TP + FP)$. *Recall* is the probability that PROACT will correctly

infer a given true activity: $TP/(TP + FN)$. The medical community calls precision and recall *positive predictive value* and *sensitivity*, respectively.

Discussion

PROACT correctly inferred that an activity occurred 88 percent of the time. For eight ADLs, there were no false positives. Of the activity instances that actually happened, PROACT detected 73 percent correctly. Given the ambiguous and overlapping activity definitions—we didn’t help PROACT make subtle distinctions between ADLs such as personal appearance, oral hygiene, and washing up—PROACT did well.

This is the first time a system has inferred nine of the ADLs. For four more (meal preparation, toileting, heating control, and medication taking), this is the first time that researchers have reported any quantitative results. For hand washing, the one ADL with previous quantitative results, PROACT’s 100 percent precision compares well to the 95 percent that Alex Mihailidis, Geoff Fernie, and Joseph Barbenel reported,⁶ but our 33 percent recall is far below their 84 percent.

The hand-washing performance is actually our worst case. Water and metal absorb the radio waves that most RFID tags use; metal can also short-circuit the tag antenna. These factors cause the detection rate to plummet for tags that are too close to those substances. This especially affected the tags on the faucets, soap bottle, and refrigerator handle. This directly affected washing hands (ADL 4), making a snack (9), or preparing a drink (10). More careful tag positioning, and using newer RF tags that are optimized for placement near liquids and metal, could mitigate this.

Sometimes, the model involves so few observations as to not be easily distin-

guishable from noise. For example, the only observable for adjusting heat (ADL 7) was the act of touching the thermostat; the only observable for playing music (11) was the act of touching the stereo. Often, a single observation of a tag provided insufficient evidence for the inference engine to confidently conclude that a new task had begun. Adding more relevant tags could solve this (for example, tagging the CDs).

Activities with starting points posed a more subtle problem. For instance, activities for personal appearance (ADL 1), oral hygiene (2), toileting (3), and washing (4) all begin with entering the bathroom and possibly turning on the light.

Our models replicate the nodes for these subactivities for each of the four activities. When PROACT detects someone entering the bathroom and turning on the light, each of these four activities is equally likely. When the system detects disambiguating objects (such as a toothbrush and toothpaste), the inference engine concludes that oral hygiene is the correct activity. If the user then uses the toilet without leaving the bathroom first, the inference engine fails to detect the second activity. To solve this problem, we might need to consider more nuanced representations of activities.

PROACT can learn parameters without supervision (in particular, duration of subactivities), similarly to the research of Donald Patterson and his colleagues.³ However, our experiments didn't include any model learning. We used a fixed small number (20 seconds) for mean duration, which was sometimes inappropriate. A particularly insidious effect surfaced when subjects interleaved boiling water for tea or coffee (ADL 10) with other activities. Because boiling takes several minutes, PROACT would conclude that they were no longer boiling water after some seconds (yielding false negatives) and erroneously jump to other activities for the post-boiling events (yielding false positives). We can achieve much better tuning of the time parameter by using unlabelled training data.

PROACT's emphasis on user interactions with objects potentially offers a powerful way to infer ADLs. As our results show, the early results are promising.

For future work in the quantitative domain, a top priority is to perform a true longitudinal study with actual elders as soon as the sensing technology permits. In the sensor domain, we want to move beyond the sensor glove and inte-

grate other types of sensors, particularly location sensors. In data mining, we're investigating how to mine better models, integrate the time of day into activity probabilities, and suggest the key objects to be tagged for a given activity set. Key challenges in activity inference include modeling interleaved activities, multiperson activities, and more complicated temporal models. Finally, we want to demonstrate PROACT's breadth by applying it to a non-ADL domain. ■

ACKNOWLEDGMENTS

We thank our subjects for their time and feedback, the Intel Proactive Health Strategic Research Project for suggesting ADL monitoring as a task domain, and the family members of the augmented house for opening their house to strangers for six weeks.

REFERENCES

1. A. McDonald and J. Curtis, *Not Alone: A Good Practice Guide to Working with People with Dementia in Sheltered Housing*, Anchor Trust, 2001.
2. A. Schmidt, H.-W. Gellersen, and C. Merz, "Enabling Implicit Human Computer Interaction: A Wearable RFID-Tag Reader," *Proc. 4th Int'l Symp. Wearable Computers (ISWC 00)*, IEEE CS Press, 2000, pp. 193–194.
3. D. Patterson et al., "Inferring High-Level Behavior from Low-Level Sensors," *Proc. 5th Int'l Conf. Ubiquitous Computing (UbiComp 2003)*, LNCS 2864, Springer-Verlag, 2003, pp. 73–89.
4. M. Perkowitz et al., "Mining Models of Human Activities from the Web," *Proc. 13th Int'l World Wide Web Conf. (WWW 2004)*, ACM Press, 2004, pp. 573–582.
5. F. Heylighen, "Mining Associative Meanings from the Web: From Word Disambiguation to the Global Brain," *Proc. Int'l Colloquium: Trends in Special Language & Language Technology*, Antwerpen, 2001,



Matthai Philipose is a researcher at Intel Research Seattle. His primary areas of interest are programming languages and probabilistic reasoning. He is currently working on sensors, data modeling, and statistical reasoning techniques for recognizing human activities. He has a MS in Computer Science from the University of Washington. Contact him at Intel Research Seattle, 1100 NE 45th St, Sixth Floor, Seattle, WA 98195; matthai.philipose@intel.com.



Kenneth P. Fishkin is a researcher at Intel Research Seattle and an Affiliate professor of computer science at the University of Washington. His current research centers on ways that ubicomp networks can help infer and guide human activities, with a focus on novel uses and techniques for RFID technology. He holds an MS in computer science from the University of California, Berkeley. He is a member of the ACM. Contact him at Intel Research Seattle, 1100 NE 45th St, Sixth Floor, Seattle, WA 98195; kenneth.p.fishkin@intel.com.



Mike Perkowitz is an analyst at Amazon.com, where he applies data-mining and machine learning techniques to the understanding of customer behavior. His research interests are in applying AI to user interfaces and Web applications. He has a PhD in computer science from the University of Washington. He is a member of the ACM. Contact him at 6910 Roosevelt Way NE, #124, Seattle, WA 98115; mike.perkowitz@intel.com.



Donald J. Patterson is a PhD candidate at the University of Washington. His interests lie in extending statistical learning algorithms to support ubiquitous computing. Contact him at 11341 B 17th Ave NE, Seattle, WA 98125; djp3@cs.washington.edu.



Dirk Hähnel is a PhD student in the Department of Computer Science at the University of Freiburg. His areas of interest lie in mobile robot navigation and especially in the learning of 2D and 3D maps. Contact him at Intel Research Seattle, 1100 NE 45th St, Sixth Floor, Seattle, WA 98195; haehnel@informatik.uni-freiburg.de.



Dieter Fox is an assistant professor of computer science and engineering at the University of Washington. His research focuses on probabilistic sensor interpretation, state estimation, and their application to mobile robotics and activity recognition. He has a PhD in state estimation in mobile robotics from the University of Bonn. He has received a National Science Foundation Career award and several best paper awards at major robotics and AI conferences. He's a member of the IEEE and AAAI. Contact him at the Dept. of Computer Science & Eng., Paul G. Allen Center, 185 Stevens Way, Univ. of Washington, Seattle, WA 98195; fox@cs.washington.edu.



Henry Kautz is a professor in the Department of Computer Science and Engineering at the University of Washington. He is a fellow of the American Association for Artificial Intelligence and winner of the Computers and Thought Award. Contact him at the Dept. of Computer Science & Eng., Paul G. Allen Center, 185 Stevens Way, University of Washington, Seattle, WA 98195; hkautz@cs.washington.edu.

pp. 15–44.

nology, vol. 13, 2001, pp. 23–39.

6. A. Mihailidis, G. Fernie, and J.C. Barbenel, "The Use of Artificial Intelligence in the Design of an Intelligent Cognitive Orthosis for People with Dementia," *Assistive Tech-*

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.