Automatic Health Monitoring Using Anonymous, Binary Sensors

Daniel Wilson & Chris Atkeson

Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213
dan.wilson@cs.cmu.edu, cga@cs.cmu.edu

1 Ubiquitous Computing in the Home

Elderly individuals living independently (often in rural areas) face challenges staying connected to friends, relatives, and caretakers. Automatic health monitoring supports the goal of ”aging in place”, or keeping elders independent and not institutionalized by providing crucial information to key individuals. Automatic monitoring can improve the accuracy of pharmacologic interventions, track illness progression, and lower caretaker stress levels [2]. Additionally, carefully filtered day-to-day information could link elders to important people in their lives.

The basic goals of ubiquitous computing are aligned with the needs of automatic health monitoring. They include identifying people, tracking people as they move, and knowing what activities people are engaged in. More challenging goals include recognizing when people deviate from regular patterns of behavior and predicting future behavior. This research uses machine learning and ubiquitous computing for simultaneous tracking and activity recognition. Identity is collected once, upon entry and exit of the home, via a single radio frequency identification (RFID) reader. Elsewhere, tracking is performed with a multitude of anonymous, binary sensors. We envision a system that unobtrusively collects information vital to cement ties between independent elders and family, friends, and caretakers. See the following scenario:

A man has an elderly mother living alone one hour away. Last week she knocked the phone off the hook and was unavailable for an entire day. The man walks into a hardware store and emerges with a large brown box. It contains several dozen nondescript, quarter-sized sensors that stick to any surface. Following directions, the man attaches sensors to doors, drawers, and chairs. He pulls out a CD-ROM and installs software on a personal computer and plugs a device into a USB port. The software instructs him to perform a quick walk-through of the house, touching every sensor. Later that week the man logs onto the Internet, types a password, and checks to see that his Mother has eaten lunch. One week later he checks that she has been cooking and eating meals. One month later he checks that her activity levels are steady.

2 Simultaneous Tracking and Activity Recognition

We wish to provide room-level tracking and recognition of ambulation. There are two main problems, (1) which person is which, and (2) where are the people? First, occupant identity is estimated and anonymous observations are assigned to the most likely occupants. This problem, called *data association*, can become severe. In an environment with m occupants and k sensors, there are m^k possible assignments. There could easily be hundreds of cheap sensors monitoring several occupants, resulting in too many data assignments to enumerate (e.g., 5^{100}). Second, these sensor readings (called observations) are used to update the location of each occupant.

Uncertainty occurs when several occupants share the same room and trigger the same set of anonymous sensors. The tracker does not know which occupant triggered which sensor (i.e., which data to associate with which occupant). There are several ways to reduce this ambiguity. The simplest is to **increase the number of ID sensors**, but this violates necessary sensor properties. Alternately, we could **increase the sensor granularity**. Higher granularity lessens the probability that multiple occupants will share the same anonymous sensor. For example, [4] uses high granularity laser range finders. We place binary sensors to maximize granularity. In a long term installation, we can **learn individual movement patterns.** Over time, motion models represent particular habits of select individuals. Also, we can **recognize activities.** Activity recognition is performed alongside tracking. Current work models ambulation, which can separate seated and moving occupants sharing the same space.

2.1 Anonymous, Binary Sensors

There are severe constraints on sensors used for long term home use by the elderly. They should be invisible or familiar and offer zero impact to daily routines. They should be cheap and readily available, as well as easy to install, maintain, and replace. They should collect only enough information to solve the task at hand, preserving privacy and minimizing computational load. Although we explore vision and auditory systems to do these tasks [1], we find that *anonymous* and *binary* sensors satisfy many of these properties. Anonymous sensors satisfy privacy constraints because they do not directly identify the person being sensed. Binary sensors, which report a value of zero or one at each time step, satisfy computational constraints. These sensors exist in many home security systems and have become a familiar sight. They are valuable to the home security industry because they are economical, easy to install, require minimal maintenance and supervision, and are robust to damage. We choose them for the same reasons, and because they *already exist* in many of our target environments (although we typically use a denser installation). Specifically, we use **motion detectors**, **contact switches**, **break-beam sensors**, and **pressure mats**.

2.2 Particle Filter Approach

Bayes filters estimate the state of a dynamic system from noisy sensor data in real world domains. The *state* represents occupant location and activity, while sensors provide information about the state. A posterior probability distribution, called the *belief*, describes the probability that the occupant is in each state. A Bayes filter updates the posterior probability density over the state space at each time step, conditioned on the data. They model systems over time through the Markov assumption that the current state depends only on the previous state. For simplicity, we assume independence between occupants, although there are joint solutions (e.g., [6] track multiple interacting ants).

Particle filters approximate the Bayes filter and solve tracking and data association problems. The key idea is to sample possible states and data associations and choose the one with maximum likelihood. The particle filter maintains a variety of hypotheses regarding the identity of each occupant, with each particle representing a hypothesis of occupant identities, locations, and activities. The data association problem can quickly become intractable as more occupants and sensors are added, so the approximation provided by particle filter methods becomes crucial. Specifically, we use sequential importance sampling with re-sampling, which approximates the Bayes filter update using a sample based representation [3].

Individual behavior is represented by unique transition probabilities between rooms and activities. The easiest way to learn these parameters is to use information gathered when only a single occupant is home (as indicated by the RFID sensor). This method ignores a significant amount of training data because occupants are often home together. In this situation, a common method to minimize uncertainty is to use the Expectation-Maximization (EM) algorithm. A version of the EM algorithm called Monte Carlo EM [8] takes advantage of the set of particles representing the posterior. A similar approach was used by [7] to learn the position of a traveler using GPS readings.

2.3 Experimental Results

We measured tracker performance in the real world over a five day period for one to three occupants. The house contained one RFID reader, twenty four motion detectors, and twenty four contact switches. The tracker used individual motion models for the three occupants. The tracker correctly classified 84.6% of the experiment. Accuracy drops to 73.7% when sleeping periods are removed.

3 Activity Labeling - Augmented Recall Survey

We plan to expand our activity recognition into several Activities of Daily Living (ADLs). Recognizing a specific activity requires examples of that activity occurring. There are many approaches to labeling this data, including interviews, direct observation, self reporting, off-line review, and experience sampling methods (ESMs) [5]. We explore a modified off-line review system, called the *Augmented Recall Survey* (ARS).

First, rough examples of activities, called *episodes* are automatically segmented from the data on a room by room basis. We ignore periods of time with more than one occupant. Second, similar episodes are automatically clustered into groups. A prototypical episode from each cluster is chosen. We use a service called the Narrator [9] to convert the episode to English text. The Narrator is a finite state machine that parses movement and activity information provided by a people tracker and generates a concise, readable summary. These summaries are presented to users later for hand labeling. We are currently preparing a user study to determine the effectiveness of this approach.

4 Expectations for the Workshop

The workshop will provide the first author with the opportunity to meet other researchers working on similar problems with similar technologies. It offers a chance to achieve consensus among researchers as to problem areas in need of further research.

Specifically, it is interesting to gauge exactly what information researchers find valuable. How are other researchers gaining access to spaces (e.g., homes, laboratories, etc.,)? Which sensors are they using and what subsequent issues are they facing? We recognize that our sensor choice represents one of many potential configurations, each with its own merits and drawbacks.

5 Biographies

Daniel Wilson is an NSF Graduate Fellow pursuing a Ph.D. in the Robotics Institute of CMU. His research goal is to provide low cost tracking and activity recognition via many anonymous, binary sensors. He is also investigating the utility of computer generated English summaries for hand labeling of everyday activities.

Chris Atkeson is a CMU Associate Professor in Robotics and Human Computer Interaction. His previous projects include Georgia Tech's Aware Home and Classroom 2000.

References

1. S. Allin, A. Barucha, J. Zimmerman, D. Wilson, M. Roberson, S. Stevens, H. Watclar, and C. Atkeson. Toward the automatic assessment of behavioral disturbances of dementia. In *UbiHealth 2003: The 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications. Seattle, USA*, 2003.
2. L. Davis, K. Buckwalter, and L. Burgio. Measuring problem behaviors in dementia: Developing a methodological agenda. *Adv Nurs. Sci.*, 20(1):40–55, 1997.
3. A. Doucet, N. de Freitas, and N. Gordon. In *Sequential Monte Carlo in Practice*, 2001.
4. D. Fox, J. Hightower, L. Liao, D. Schulz, and G. Borriello. Bayesian filtering for location estimation. *IEEE Pervasive Computing*, 2(3):24–33, July-September 2003.
5. S. Intille, E. M. Tapia, J. Rondoni, J. Beaudin, C. Kukla, S. Agarwal, and L. Bao. Tools for studying behavior and technology in natural settings. In *Proceedings of UBICOMP 2003*, 2003.
6. Z. Khan, T. Balch, and F. Dellaert. Efficient particle filter-based tracking of multiple interacting targets using an mrf-based motion model. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'03)*, 2003.
7. D. Patterson, L. Liao, D. Fox, and H. Kautz. Inferring high-level behavior from low-level sensors. In *UBICOMP 2003: The Fifth International Conference on Ubiquitous Computing (October 12-15, 2003)*, 2003.
8. G. Wei and M. Tanner. A monte carlo implementation of the em algorithm and the poor man's data augmentation algorithms. *Journal fo the American Statistical Association*, 85(411):699–704, 1990.
9. D. Wilson and C. Atkeson. The narrator : A daily activity summarizer using simple sensors in an instrumented environment. In *The Fifth International Conference on Ubiquitous Computing (UBICOMP) 2003 Demonstrations, Seattle, USA*, 2003.