

The Aware Home: A living laboratory for technologies for successful aging

Gregory D. Abowd, Aaron F. Bobick, Irfan A. Essa, Elizabeth D. Mynatt, Wendy A. Rogers

GVU Center Georgia Institute of Technology
Atlanta, GA 30332-0280, USA

Abstract

We describe our ongoing research in the area of developing and testing technologies for successful aging. Sensing and perception technologies can enable a home environment to be aware of the whereabouts and activities of its occupants. Motivated by the desire to use such an awareness to help maintain independence and quality of life for an aging population, we describe the technological, design and engineering research challenges inherent in this problem domain. Our work is situated in the Georgia Tech Broadband Institute's Residential Laboratory, a unique living laboratory for this exploration of ubiquitous computing in a domestic setting.

Introduction

How can your house help you and your family if it is made aware of your whereabouts and activities? To answer this question, and other questions related to the implications of ubiquitous computing technologies in the home, Georgia Tech opened the doors on May 1, 2000 a new research facility, the Broadband Institute's Residential Laboratory (see Figure 1). This 5,040 square foot laboratory — a two story house containing two identical, two bedroom apartments, basement and attic — has been specially designed and constructed to facilitate advanced research and development as well as extended human subjects experimentation.

We have initiated a long-term research program, the Aware Home Research Initiative (AHRI), to push on the specific multidisciplinary challenges of building technology and appropriate services that leverage awareness of human activity in the domestic setting. In this paper, we will address the motivation and research challenges of this research initiative. We have created a unique living laboratory for investigating future

ubiquitous computing technologies in a domestic setting. There are significant technological and engineering challenges to overcome in order to create an environment that can perceive the activities of its occupants. Equally significant is the challenge to invent a satisfactory human experience amidst the awareness. We will explain our human-centered motivation for promoting awareness in a domestic setting and then address some of the design, technology and software engineering challenges we are pursuing.

Motivation: Maintaining quality of life for an aging population

America's population is aging. The U.S. Census Bureau reported nearly 44 million adults over age 60 in 1996, and the projected number for 2025 is approximately 82 million (over 20% of the total population). Baby boomers approaching late middle age ask, "How does one care for a population that lives many years longer than any of the preceding generations?" Furthermore, the costs of traditional assistive care will continue to be out of reach for the majority of Americans, costing on average \$60,000 per year, and often requiring a substantial entry fee (Greenwald 1999). These costs are more prohibitive as the demographics in the U.S. continue to change. Specifically, the ratio of wage earners to senior adults will shift from 4-1 to 2-1 in the next 25 years (Vanderheiden 2000).

For economic reasons alone, increasing the length of time that individuals can defer assistive care is valuable. For reasons beyond economic ones, a primary goal of many older individuals is to maintain an independent lifestyle. Thus, many older adults live in private homes, typically either alone or with family. Our own ethnographic studies of local Atlanta assistive care centers focused on understanding why people chose to leave their homes for an institutional care setting. We were impressed by the significant and sometimes dominant role of adult children and their need for day-to-day assurance of their parent's well being. We observed the growing emphasis on cognitive declines.

We also noted the negative effects of institutional life, for example, decreased motivation for self-care, and increased isolation and depression.

Based upon these observations, we believe that the development and careful placement of appropriate technological support can empower older adults to continue living in their own homes longer. We specifically focus on the effects of declining cognitive abilities on independent living, and the role information technology can play in augmenting those capabilities. We have identified two areas that are key to sustaining quality of life for an independent senior adult:

Maintaining a Familial Vigilance A critical function of an assistive care setting is to provide a safe, healthy environment. This clearly requires detecting crises or emergencies. A more demanding aspect of providing such an environment, however, is the more general monitoring of the occupants' situation and the noting of significant changes in the manner in which they behave or interact. We refer to such extended monitoring over time as *trending*. Examples include noticing reduced mobility or far less time engaged in cognitively stimulating activities. Furthermore, the information about an older individual should be shared in a privacy respecting manner, sharing appropriately filtered information with concerned family members and trusted caregivers, thus promoting a peace of mind that comes with mutual awareness and day-to-day assurance.

Supporting Daily Routines: Declines in cognitive and sensory capabilities manifest themselves as difficulties in performing many daily tasks, such as taking medication, preparing food, and operating household appliances. An aware environment should help offset typical cognitive impairments, such as problems with retrospective memory — the episodic remembering of what the occupant was doing recently, say, before being interrupted. Providing relevant information at the appropriate time (e.g., which medications should be taken now that the occupant has awakened and is headed to the bathroom), reduces the cognitive burden of doing the everyday activities essential to independent living.

Outlining of the research questions

Though human needs motivate the research in the Aware Home, two principal technical intellectual challenges are being explored simultaneously:

Human-Computer Interaction The design of devices, services and interfaces to facilitate novel human-human and human-environment interactions



Figure 1: *The Georgia Tech Broadband Institute Residential Laboratory, home for the Aware Home Research Initiative. (a) The floor plan for one of two identical living floors. (b) The exterior view of the property.*

in light of increased awareness, with particular focus on interactions for an aging population and their distributed families.

Computational Perception The development of techniques with a variety of sensing devices that lead to an environmental awareness of human activity in the home.

Designing novel interactions in the Aware Home

Research in Human-Computer Interaction encompasses examining the usefulness and usability of future technologies. Both of these goals are particularly challenging in home environments. In the context of this proposal, usefulness focuses on supporting independence and quality of life. Much of our research is guided by understanding typical cognitive declines that accompany aging and their implication for independent life. Moreover communication technologies can improve the quality of life by increasing social networks, daily interaction and feelings of connectedness. To support independence, usability focuses on providing support while not undermining the person's ability to care for themselves. For seniors, it is important to design interfaces that take into account many of the typical declines in perception and cognition.

Supporting familial vigilance and communication

In addition to supporting the needs of older adults, we have discovered that it is critical to support communication and awareness between extended family members (Mynatt, Essa, & Rogers 2000). An aging adult's desire to remain in the familiar setting of the family home frequently must be balanced with their extended family's desire to keep them safe. Clearly this balance becomes more precarious as age increases. An aging couple can support one another, but if one becomes incapacitated, can the other support himself? No longer

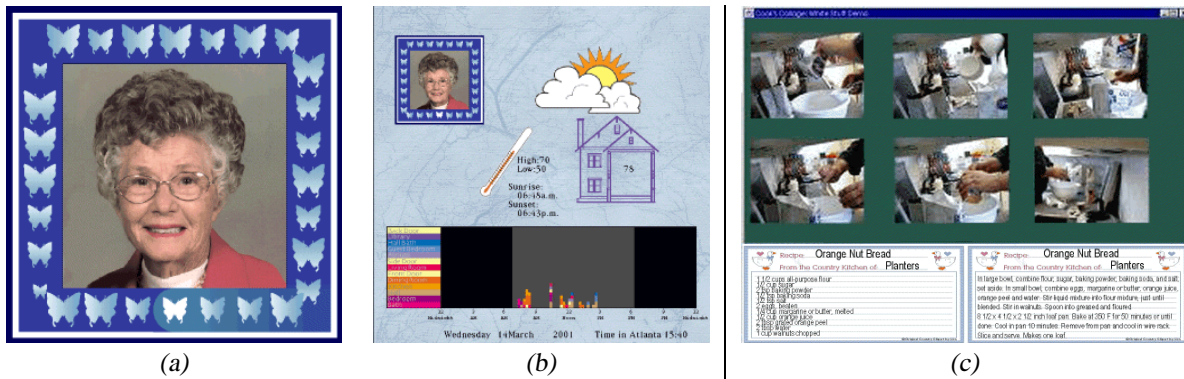


Figure 2: The Digital Family Portrait. (a) illustrates 28 days of activity information about the woman pictured. (b) shows detailed information about one day including the weather, indoor and outdoor temperature, and the number of room-to-room transitions in 15 minute increments. (c) A surrogate memory display. Snapshots of six recent actions are shown in order (upper left to lower right) for a cooking activity.

can family members have the peace of mind that derives from knowing that one aging parent can support the other. Additionally geographic distance between extended family members exacerbates the problem by denying the casual daily contact that naturally occurs when families are co-located. Providing support for light-weight communication and remaining aware of a distant family members' day-to-day activities, what could be called surrogate social support, can promote the peace of mind necessary for those senior family members to age in place.

Designing a Digital Family Portrait: Our digital family portrait addresses the “peace of mind” deficit that often compels adult children to move their aging parent to an institutional care setting (Mynatt *et al.* 2001). Our goal is to support awareness of the long-term health, activity, and social well-being of senior adults living by themselves, answering questions such as, “Has she been eating enough?” and, “Is he active or sedentary?” Like a traditional portrait, it is designed to be hung on the wall or propped on a mantle, blending with household decorations. In our interface, we represent the well-being of family members with iconic imagery, based on sensing data from their home. From general measurements of activity, to indications of the weather, the portrait attempts to capture the observations that would naturally occur to someone living in the same home or next door. Most awareness interfaces only provide a snapshot of the present. Since many questions about an elderly parent refer to trends over time, such as, “Is she becoming more socially isolated?,” we provide representations of the past, as well as the present.

Our current design of the digital family portrait is

shown in Figure 2b. Each icon represents a day; the current day is white with time moving clockwise. In this portrait the size of the butterfly indicates how active the person has been in his home. By touching the butterfly, the viewer gets information about that day. In this example, activity is based on room to room movements in 15 minute increments. A graph at the bottom of the screen indicates room transitions throughout the day. Information about current weather and temperature conditions at the remote location are also included.

Our initial research indicates the feasibility of creating a visual representation of a senior adult’s well-being and sharing that information with select family members in their home (Mynatt *et al.* 2001).

Providing serendipitous information in support of daily activities: An oft-cited frustration of daily life is not having the information you need when you need it. However, as people age, this missing information is less likely to be a current stock price, and more often one of the details in a common routine. For example, as retrospective memory capabilities decline, a common memory lapse is failing to recall recent events in the midst of an interrupted task (Rogers *et al.* 1998). Difficulties in successfully completing household tasks, such as preparing a meal, can undermine a person’s self-confidence, and hence their ability to live independently. Likewise, both cognitive and perceptual declines may lead to difficulties in recognizing changes in the physical environment, precipitating a fall or the inability to recover a moved object (Smith 1990). In this research, we investigate the design of interfaces that provide a constant, but unobtrusive, flow of information that may be used serendipitously in the midst of common daily activities.

Surrogate Memory Systems: In our *What Was I Cooking?* interface, we are investigating how sensing and capture capabilities can support recovery from common memory lapses that occur when a person is distracted while cooking. As cooks perform actions in our kitchen, sensors register these events and trigger the grabbing of snapshots from a continuously running video stream. These snapshots are currently arranged using a cartoon panel metaphor; six to eight images in two rows (see Figure 2a). Questions that our interface may help people answer are, “Did I add the baking powder?” and, “When did I put the roast in the oven?” Our goal is for users, by glancing at our visual displays, to ascertain the answers to those questions with minimal effort. In order to further design and evaluate this interface we are conduct a series of Wizard of Oz empirical studies.

One common memory lapse is referred to as “loss of activation.” Specifically, a person will initiate a task in a certain context (e.g. glancing at a calendar and remembering to pay a bill), and then move to a new context and forget the remainder of the task (e.g. left standing in the study with no memory of why he went there). In these memory lapses, the “interruption” is self-initiated as the person shifts context. A recovery technique is for a person to retrace his footsteps until he recovers the original context and catalyst for the task. Similar to the *What Was I Cooking?* interface, a room-by-room record can aid in bringing that context with the person as they change rooms.

Making the environment aware of human activity

The novel interfaces just described demand a range of perceptual capabilities to uncover the context of everyday life in the home. The computational perception research focuses both on providing base-line tracking performance (*who* is *where* and *when*) and on developing fundamental new methods for determining qualitative behavior (*what*).

Location and orientation: A significant number of potential applications require knowledge of the location of the occupant or occupants at all times. One of the major categories in the ADL description of routines is that of “ambulation” — a measure of a person’s ability to get around in their environment. A rudimentary assessment of a persons mobility can be established by merely knowing the location of that individual at all (or most) times.

As an initial attempt to provide location information quickly in the Aware Home, we have custom-designed an RF ID system that uses antennas in the form of floor mats distributed throughout the house. Individuals wearing passive RF ID tags below the knee can

now be tracked by simply recording the traversal over the strategically placed mats (see Figure 3). This system works now and can provide room-level occupancy information of known individuals in the house. What we will describe now are more sophisticated methods to obtain location and orientation information throughout the home.

We have installed an unobtrusive camera grid on the first floor in the Aware Home. These *optical sensors* are installed in the ceiling tiles at a height of 3 meters, with a FOV of approximately 120 degrees; the camera and a sample image are seen in Figure 4. These cameras are calibrated with respect to the location in the house. We have also developed a client/server “blob” management system where sensor clients perform background subtraction on a per camera basis and report detailed information of each significant foreground connected component to the server process. The server process maintains a spatially and temporally addressable database of these blobs that consumer clients can access. We have already implemented some simple gesture recognition clients to act as a triggers from a person to the environment to initiate some further interaction, such as voice recognition.

We are implementing a “track-stitching” algorithm (similar to our real time closed-world tracking approach in (Intille, Davis, & Bobick 1997)) that will give temporally-extended tracking information per person. One of the benefits of conducting this research in the situated context of the home is that the interface applications establish performance criteria. For example, if there are two people in the house and they sit together on a couch for a period of time, the system may have trouble determining who is who, especially from an overhead view. However, as long as once the two people separate the system can regain identity, then the confusion during the couch interval is unlikely to affect performance in tasks such as assessing mobility. Furthermore, since assessment does not require causal processing, it is possible to use future information to disambiguate uncertainty.

Even given accurate tracking information there is work necessary to correlate properties computable from time-stamped tracking data with ADL assessments. We expect that the tracking capability will provide the earliest technology-based collaboration between technology developers and those researchers focused on assessing the quality of life of an aging population. Once tracking is in place, we will perform initial experiments in which independent observers provide Ambulation ADL labelings about subjects that are tracked by the vision system. These labelings will allow us to apply standard pattern recognition techniques to the properties of the trajectories from which we need to generate such assessments. Furthermore, any such

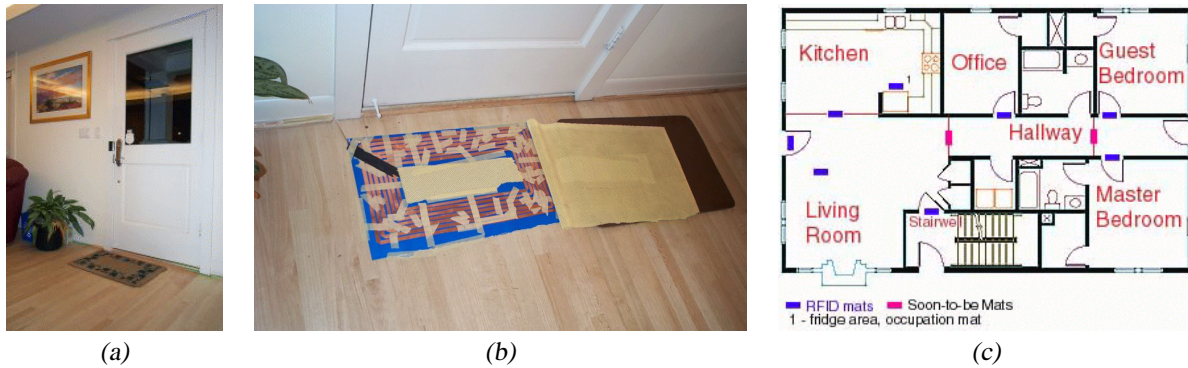


Figure 3: The RF ID room location system in the Aware Home. (a) The RF ID antenna are hidden under doormats, as shown, or under carpets within the home. (b) A view of the hand-crafted RF ID antenna. (c) A floorplan indicating positioning of antenna/mats throughout one floor of the Aware Home

automatic labeling, regardless of its correlation with human assessment, will be able to provide trend analysis of a given occupant over extended periods of time.

There is also fundamental work to be done in determining the orientation and/or pose of a person. Such a description is necessary for, say, the *context-establishing capture* required by the “room-to-room” reminder system discussed above. To collect an image that will serve to remind an occupant what they were *intending* to do when they entered a particular room, we will need to capture an image in the previous room that shows either what they were doing or at least to what they were attending. While the design research determines the most appropriate image to display, the perception research must be able to determine the relevant parameters.

We are pursuing two separate approaches to pose/orientation estimation. The first method is strictly appearance-based. The fundamental idea is to add temporal continuity to a silhouette-driven, body-part labeling method. Our second approach is to take advantage of the multiple oblique views available in each room of the house. We make use of an image-based visual hull (Matusik *et al.* 2000) to construct the approximate three-dimensional volume of the person from a set of calibrated cameras or a calibrated environment. For any such volume there is a limited range of body poses consistent with the approximate model. We are also pursuing an approach to detect features on the visual hull that can aid in tracking of limbs over time to refine the skeletal shape, with its inherent articulatory constraints.

Local behavior recognition: Though much contextually relevant information can be gleaned from the geometric information of location and pose, it is clear that deeper inferences can be drawn if some informa-

tion is provided about which activity is taking place. Consider a simple example — the location and tracking system can determine the occupant is sitting on the couch. But, there is a significant difference from a cognitive engagement perspective between sitting on a couch watching TV (typically cognitively passive) or reading a newspaper (more likely to be cognitively engaging). Determining in which activity the resident is engaged increases the sophistication and efficacy of the familial vigilance.

The reading scenario is an example of simple *local* behavior recognition. By local we mean that the behavior itself can be defined and therefore recognized in terms of elements that are local in terms of time and space. For example, the presence of the newspaper, the location on the couch, and a repeated low duty-cycle lateral motion is probably sufficient to recognize newspaper reading. There is no extended interaction between the occupant and the multiple objects.

Most previous work on human activity recognition, especially at the level of say gesture or simple human motions, has focused on this type of problem. A common element of these approaches is that are probabilistic, assigning a likelihood that an observed sequence of features is generated by a particular behavior. In our work here we need to include contextual information in the local behavior recognition methods. For example, simple attributes such as location or time of day can be exploited to yield much more robust and effective recognition procedures.

A more subtle form of contextual influence is used to decide between competing behaviors. Again, suppose the person is sitting on the couch. The system maybe uncertain whether the newspaper is present, and furthermore the motion observed could be page turning, or channel surfing with a remote control. The simple contextual fact that the television is not on is enough to

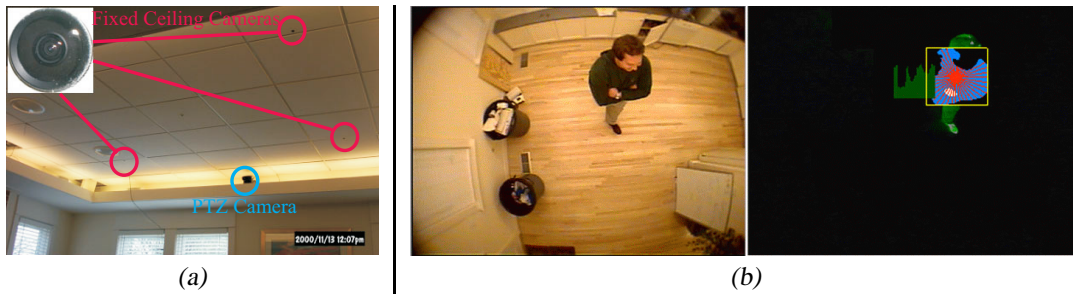


Figure 4: (a) Living room in the house, showing several small cameras installed in the ceiling tiles. Also shown is a pan-tilt-zoom (PTZ) camera. (b) View from the ceiling cameras in the kitchen, with segmented person.

reduce the prior probability of channel surfing to zero. Thus, the context about the television affects the system’s willingness to believe that the occupant is reading. This is an example of the conditional dependencies that arise in probabilistic reasoning.

Extended behavior recognition: The final and most ambitious effort in the computational perception area is the recognition of *extended behavior*. By extended we mean either in time or in time and space. For example, the act of preparing a meal is not restricted to a narrow location, and is best determined by the *interaction* between the occupant and various other objects and locations in the scene. To recognize such activities requires a representation that makes explicit the temporal and spatial relationships between component elements. For example, the moving of objects from refrigerator to counter, counter to table, table to dishwasher, with an indeterminate number of stops back to the refrigerator is highly indicative of someone preparing a meal and then eating it.

To perform the recognition of these types of activities, we apply two related yet structurally distinct methods. The first of these is a structural analysis approach that uses a *stochastic context-free grammar* (SCFG) to analyze incoming video frames. The attraction of such a method is the opportunity to provide a priori grammatical definitions of an activity. For example, one can define meal preparation as requiring activity near a pantry or refrigerator before sitting down at the table with a plate of food. The probabilistic nature of SCFGs permits robust performance in the presence of uncertain low-level feature detectors. It is an open question as to whether we can design low-level detectors of sufficient power to support the type of grammatical reasoning required. Determining what the appropriate features or primitives are is where much of our initial effort would be focused.

Handling extremely noisy detectors is the goal of an alternative approach. This method employs belief net-

works to probabilistically weigh perceptual evidence. The insight here is that extended behaviors can be recognized by observing relatively well structured local events and then combining the observation of those events to determine larger scale activity. We employ exactly the types of detectors designed for the grammatical parsing approach above, but apply them in the more forgiving belief network integration method.

Conclusions

We have described the long-term multidisciplinary research agenda of the Aware Home Research Initiative. Situated in a unique living laboratory, we are in a very good position to explore the technical and human implications of ubicomp in a domestic setting. From the human perspective, our research is driven currently by the main theme of aging in place. As researchers, we see significant intellectual challenges in the domains of Computational Perception, Human-Computer Interaction and Software Engineering and we have used this paper to review our progress and future goals in these distinct areas.

Any significant research problem covers many more challenges than we have been able to address in this brief overview. We realize that much of our work in designing applications for an aware environment can be generalized to cover more universal design issues. Indeed, several applications projects not reported here focus on other demographics in the home, such as families with small children (Stevens, Vollmer, & Abowd 2002). Other research efforts are rightly focusing on more general health issues (see University of Rochester’s Center for Future Health at <http://www.futurehealth.rochester.edu>) or issues of designing the physical space to accommodate new technologies (see MIT’s Changing Places/House_n project at http://architecture.mit.edu/house_n).

Though what we presented in this overview does not directly address social implications of our work, we are explicitly examining socio-legal implications of sens-

ing in a home environment through a structured exploration of privacy litigation. This work is leading toward a defined process to help technologists and designers discover how the specific details of sensing and the use of context will be perceived in the wider social and legal contexts of modern life.

References

Greenwald, J. 1999. Elder care: Making the right choice. *Time Magazine* 52–56.

Intille, S.; Davis, J.; and Bobick, A. 1997. Real-time closed-world tracking. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 697–703. IEEE Computer Society Press.

Matusik, W.; Buehler, C.; Raskar, R.; Gortler, S. J.; and McMillan, L. 2000. Image-based visual hulls. In Akeley, K., ed., *Siggraph 2000, Computer Graphics Proceedings, Annual Conference Series*, 369–374. ACM Press / ACM SIGGRAPH / Addison Wesley Longman.

Mynatt, E. D.; Rowan, J.; Jacobs, A.; and Craighill, S. 2001. Digital family portraits: Supporting peace of mind for extended family members. In *Proceedings of the ACM SIGCHI 2001 Conference*, 333–340. ACM Press.

Mynatt, E. D.; Essa, I.; and Rogers, W. A. 2000. Increasing the opportunities for aging in place. In *Proceedings of the 2000 Conference on Universal Usability (CUU 2000)*.

Rogers, W. A.; Meyer, B.; Walker, N.; and Fisk, A. 1998. Functional limitations to daily living tasks in the aged: A cocus group analysis. *Human Factors* 40:111–125.

Smith, D. 1990. Human factors and aging: An overview of research needs and application opportunities. *Human Factors* 32:509–526.

Stevens, M.; Vollmer, F.; and Abowd, G. D. 2002. The living memory box: Function, form and user-centered design. In *Companion Proceedings of the 2002 conference on Human Factors in Computing Systems — CHI 2002*. Interactive Poster. To appear.

Vanderheiden, G. 2000. Fundamental principles and priority setting for universal usability. In *Proceedings of the 2000 Conference on Universal Usability (CUU 00)*.