Physical, Social, and Experiential Knowledge in Pervasive Computing Environments

A long-term deployment of a system for recording experiences in informal spaces demonstrates that people use physical, social, and experiential knowledge to determine new technologies' relative utility and safety.

Gillian R. Hayes, Erika Shehan Poole, Giovanni Iachello, Shwetak N. Patel, Andrea Grimes, Gregory D. Abowd, and Khai N. Truong
Georgia Institute of Technology

Pervasive computing designers and researchers often create services and applications to help people record their experiences. At the same time, cheap, small, and easy-to-deploy recording technologies are quickly emerging throughout public spaces. In many ways, these technologies are pervasive computing realized. Understanding how people deal with audio and video recording is therefore a good way to explore how people might adopt, adapt, and react to pervasive computing technologies in general.

Selective archiving is a method for recording data in an environment in which the recording devices (for example, video cameras and microphones) are always on and available. The devices record automatically into a short cache of data. If no explicit action is taken, the recording system throws the data away when it becomes older than a set time. However, a user wanting to archive some data from this cache can retrieve it retroactively before that set time.

BufferWare is a selective-archiving application we deployed in an informal space and studied for over a year. This extended study helped us understand three types of knowledge people use to form impressions of new technologies:

- *Physical knowledge* relates to a particular piece of technology's affordances—that is, the design elements that inform people about the technology and how to use it.[1]
- *Social knowledge* is a community's embedded knowledge of, for example, the conventions, laws, and uses of its space. It involves the set of social interactions occurring in or around a space housing pervasive computing technologies. Social knowledge also includes elements of trust and understanding—or suspicion and confusion—gathered from other social actors in the environment.
- *Experiential knowledge* includes the range of information from past experiences, both in new spaces and with similar spaces and technologies. These experiences let potential users and stakeholders learn about their worlds and build models into which the new experiences fit.

Here, we aim to add significantly to the research surrounding security and privacy concerns by focusing on them rather than just noting them as a side effect of testing an application's utility and usability.

The BufferWare project

Open and casual spaces present particular challenges to people wishing to record experiences:

- People might not be able to predict when events of interest will occur and often aren't prepared for manual recording.

Figure 1. The public area used for BufferWare deployment: (a) two of the tables in the room (we instrumented the table nearest the window), (b) a view of a table and clearboard, and (c) the open stairwell.

- To provide security and privacy, the recording technologies must take into account technical-infrastructure requirements and social and cultural requirements such as notification and consent.
- The need to record can be infrequent, and barriers to recording might limit the system's use.

We considered these challenges when selecting an installation place for the BufferWare system. We chose a social area next to an open stairwell on the third floor of an academic building housing media-and-design, computer science, and engineering researchers (see figure 1). The area contained three tables next to three large clearboards (large transparent writing surfaces that work similarly to whiteboards). One table was near a window. We chose this space because it was open, informal, and public—not closed off like meeting rooms, which other researchers have already studied. We instrumented the table nearest the window, so that people wouldn't have to walk through a recorded space to get to a nonrecorded area.

For six months before deployment, we recorded activity in the space using a semicontrolled manual sampling algorithm ($n > 300$ samples). The space's primary users were graduate students, staff, faculty members, and professional researchers with offices or meetings in the building. A negligible number of other people, such as undergraduate students and family and friends of workers in the building, used the space occasionally. Most of the time, the space was unoccupied, but during the afternoon, all three tables, all the chairs, and some standing space were often filled. Typical activities in this space include

- small meetings;
- lunch, snack, or coffee break gatherings;
- individual reading and work sessions;
- individuals and groups looking out the windows;
- telephone conversations;
- eavesdropping on conversations on other floors; and
- group discussions unrelated to work.

BufferWare's capture services support recording of these activity types. The system is inherently flexible, requiring minimal infrastructure (see figure 2). Users can access saved content online. A touch screen embedded in a café table provided the interface to BufferWare. A single camera, attached to a PC via video and security cables, provided video data to the application. Finally, a microphone glued to the tabletop recorded sound within a space matching the camera viewing angle.

We first deployed BufferWare from September 2005 to June 2006. We placed blue tape along the carpet to denote the recorded area, and posted signs and sent emails to alert building occupants to its presence. We also began logging requests to save or delete buffered data or review clips. In April 2006, we began using motion detectors to log motion in the space (see figure 3).

Survey and interview input from the first deployment revealed several small, correctable problems. We fixed them to help disambiguate technological concerns from ease-of-use problems. These fixes included

- enlarging buttons to accommodate people whose fingers were too wide to input easily, or who didn't have the right manual dexterity,
- expanding the buffer length to one hour for longer meetings,
- doubling the video capture resolution, and
- providing a URL with a "magic key" that automatically authenticated users to view newly saved clips.

Additional changes to the primary interface addressed concerns about visual feedback and video data access, including an expanded viewing window and notices about policies for protecting saved data. We deployed a second version of BufferWare during September and October 2006.

After the second deployment, we removed the BufferWare hardware, which included the camera, microphone, and touch-screen interface. We continued to gather motion detector data for

Figure 2. The BufferWare system: (a) A touch screen in the table provides two buttons, one for saving and one for purging the data cache. (b) Users can choose how much to save using a double slider and can send the data to registered users. (c) A Web interface lets users view and comment on saved videos.

the next 30 days to record activity in the space after the equipment had been removed (the postdeployment condition). The online service continues to allow access to saved content but not to log and analyze the interactions.

Anyone in the building could participate in the BufferWare study in various ways:

- using the service itself (36 people did this),
- completing anonymous surveys (32 people took surveys: 13 for the first deployment and 19 for the second),
- participating in interviews (we held 27 interviews with 22 people: five people after the first deployment only, 12 after the second only, and five after both).

Nearly all the interviewees also completed surveys.

We used sensors and the BufferWare software to log user interactions with or near BufferWare. We attached black tape over portions of the motion detectors to limit their view to the zone depicted in figure 3. The log associated with these motion detectors contains the date and time stamp for each sensor firing.

We examined the space's use during the first 30 days of each deployment and the first 30 days of the postdeployment phase. Across all three time periods, the space was occupied approximately 20 percent of the time. (Although the system was on only 18 hours a day, from 6 a.m. until midnight, we used a full 24-hour day in these calculations. That's because without careful inspection, people didn't necessarily know the hours of service.) People used the space approximately 18 seconds more per day during the second deployment than in the first (two-tailed matched t-test, p < 0.005) and approximately 10 minutes more per day after we removed recording from the space (two-tailed matched t-test, p < 0.001). It's hard to be certain why these differences exist. Although they're statistically significant, they're small. Furthermore, in self-reported survey data of space use, no significant differences existed between conditions. So, we would need further exploration or a replication of the study to better understand these results.

We also recorded interactions with the software through automatic logging. Across both deployments, people created 174 archives. Users accessed these video clips 257 times during the entire study (both deployments and postdeployment). Most clips were accessed. Users cleared the buffer manually 598 times by pressing the Delete button rather than waiting for the system to delete the data. Some of these presses were in a row—likely, a person trying to ensure that the deletion had registered.

Knowledge in privacy and security determinations

Here's how participants used the three types of knowledge in relationship to security and privacy while interacting with BufferWare.

Physical

Identifying appropriate ways to notify people that they're being recorded and to let them provide feedback about that recording are significant issues for researchers in security and privacy and for designers of new recording technologies. Minimizing the recording's noticeable impact in a space makes it hard for people to know what's happening and how to control it, creating usability and sociocultural issues. At the same time, designing for maximum notification about recording can be overwhelming, given the sheer volume of recording in public places, and can make users self-conscious. So, designers must strike a balance to optimize users' ability to create appropriate mental models of recording from
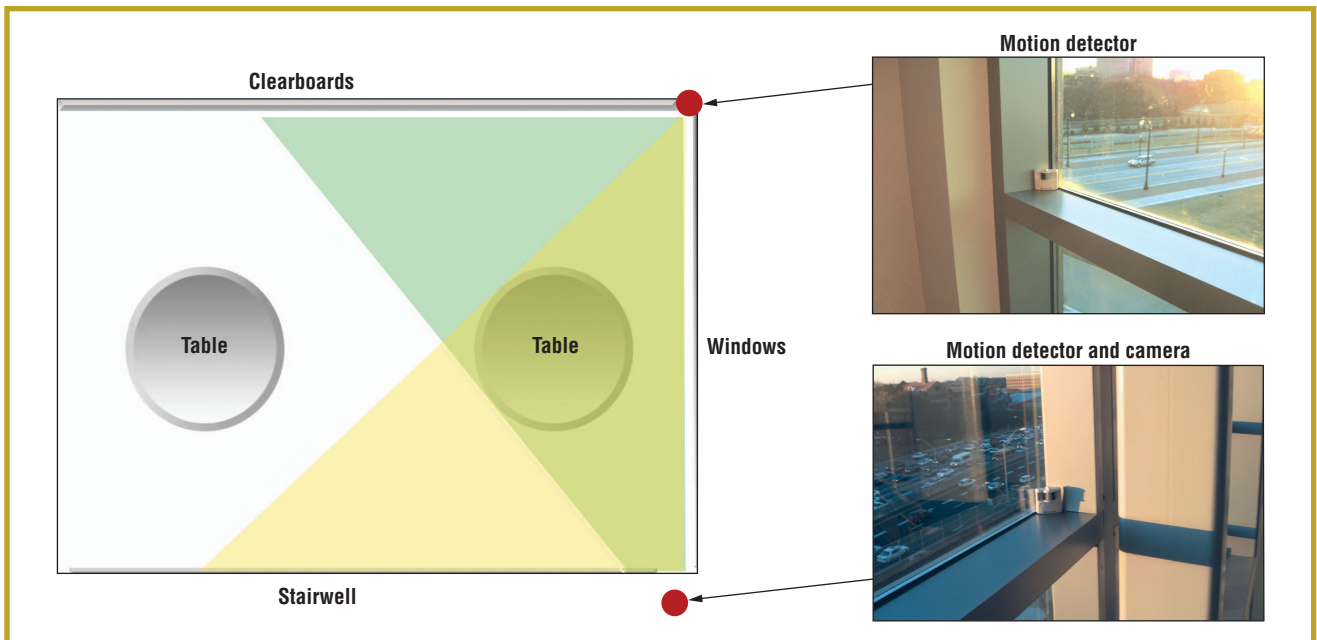
Figure 3. Two motion detectors let us track movement in the camera angle and near the table separately.

the available cues without creating other challenges.

We intentionally made visible notification and feedback in the BufferWare project's recorded space plentiful. Specifically, BufferWare's design included a visible camera and a small window, embedded in the touch screen on the table, showing the live video capture feed. As we mentioned earlier, we placed numerous signs in and outside the space and changed them substantially between the first and second deployments. As with the other issues we explored, interview and survey participants provided abundant information about physicality and visibility, including notification and feedback concerns.

People's responses to the visible camera were divided. General distaste and discomfort with cameras were some of the reasons reported for not liking Buffer-Ware. Interviewees often noted that security cameras are more acceptable because they're often inconspicuous. Some commented that they became nervous or self-conscious when having their images recorded. For example, one interviewee noted that security cameras are less of a concern, because "I'm not looking at the camera directly and … sometimes I'm not

even aware that there's a camera there. So, when I notice the camera I get nervous." On the other hand, many interviewees commented that they preferred visibility and transparency to any recording.

One interesting issue is that even a public space isn't one-dimensionally so. In Gavin Jancke and his colleagues' experiments, a vocal minority of users reported distaste for an always-on system linking two public spaces.[2] These users described many private activities that took place in the public space (personal phone calls, eating lunch, meetings, and so on), and that the private nature of this public space was disrupted.

Reflection of the capture stream is another way to provide feedback about the recording experience, a notion central to David Nguyen and Elizabeth Mynatt's "privacy mirror."[3] You could alert people who might be recorded to the presence of these technologies by using a video display of the data being captured. In BufferWare, the small screen embedded in the touch screen served this purpose. Our choice of a horizontal display rather than a vertical wall display wasn't a trivial one. We wanted to give feedback to individuals in the space

without providing too much of that information to people outside the space, who might see a vertical display.

Some individuals reported that the information was harder to view on the horizontal screen: "Maybe you need to have an LCD on the wall that's showing you 'Hey! This is what's being recorded right now!'" This same individual commented that a vertically mounted display might resemble a mirror and could contribute to a "surveillance feel" that might make a person uncomfortable or self-conscious—further demonstrating users' conflicting sentiments.

People also reported that simple physical differences between the BufferWare-enabled table and the other tables in the space added to their feelings of being part of something new, different, and possibly uncomfortable. For security reasons—both information security and physical theft—the server and display units were locked to the table. We also used hard-wired cables to transport the data instead of less secure wireless signals. These differences often had little to do with the recording itself, but still marked the table as different and somehow "not for [them]," as one participant commented.

Furthermore, usability considerations and interface choices impacted how people understood the system. The system intentionally hides complex information about data storage and deletion to make it more comprehensible. Similarly, to make

> Knowing what other people think, talking with other people affected by the system, and the general pressures of belonging to a group can all affect people's perceptions of technology.

the interface usable given a touch screen's constraints, we made certain compromises. For example, rather than ask users to type an email address to save data (as we originally designed), which more maximally protects privacy, the BufferWare interface displayed registered usernames as quick buttons. To save data, users simply pressed the buttons to select their accounts. Although some users created usernames that weren't easily identifiable, this display still presents potential risk.

Institutional considerations strongly influenced how we notified people about our recording. A legal representative for Georgia Tech's Office of Sponsored Programs prescribed the wording on the initial notification signs as well as the agreement at the time of archiving indicating that anyone in the space is comfortable with recording. During the first deployment, however, people perceived the signs as intimidating and scary. For the second deployment, we used larger, more colorful signs that stressed the technology's potential uses. The required wording remained on the signs, however. So, the signs might have continued to convey a sense of danger.

The blue tape line in the first deployment notified people about recording in the space. Many people commented that this explicit notification made them uncomfortable because it seemed to indicate

potential risk. To combat these concerns, we removed the tape during the second deployment. People interviewed after the second phase who had been in the building during the first phase almost unanimously agreed that its removal made the space seem less threatening. So, these clear indicators might increase people's awareness at the expense of a reasonable ability to judge the level of concern they should have about the recording technologies in the space.

The use of multiple indicators warning people that they might be recorded could have had an effect opposite to that intended. We thought this intentional visibility would enhance the trust level between people using the space and the researchers. However, it might have prevented potential users and the researchers from showing each other that they could be trusted through other means. Thomas Erickson and Wendy Kellogg argued that this natural trust developed over time in their Babble chat system deployment.[4]

How best to notify people about recording has been a significant challenge for BufferWare and for other capture systems. Because people differ so widely in their notification preferences (for example, some want to be notified every time a recording occurs, others wish to be notified the first time only, and still others prefer no notification), a one-size-fits-all solution is unlikely.

These experiences and concerns exemplify the conflicted responses we received when querying about physical knowledge. People want to make informed decisions about video capture, and part of

that process is knowing (and sometimes seeing) what's being captured. However, constant feedback in the physical environment can create an unnecessary aura of danger, making people more self-conscious and concerned than they might have been with less feedback.

Social

Social cues can be extremely important for building models of security, privacy, and trust in a system. Knowing what other people think, talking with other people affected by the system (or in charge of it), and the general social pressures of belonging to a group can all affect people's perceptions of technology. The BufferWare project was no exception to this pattern. During the two test deployments, people reported changing their minds about the system after seeing other people using or avoiding the space or after talking with other people about it. People often queried others with whom they were interacting when deciding whether to use the space. If someone wanted to use it, the entire group often would. If one person was uncomfortable, the entire group might acquiesce and avoid it.

We used the email distribution lists for building occupants to communicate system status changes and other information. During the second deployment, however, interview and survey responses indicated that many people, particularly those who were newest to the building when they first encountered BufferWare, didn't know what was happening in the space. After further probing, we learned that some people who had moved into the building in the past year weren't on our email lists. Instead, they reported relying on what other people were saying, a situation similar to that encountered when a commercial product is released or a government-sponsored initiative is put in place to deploy new technologies, and product buzz determines reactions.

Other people's use of BufferWare gave important social cues to potential users that it was safe in terms of security and privacy. These cues include witnessing other people in the space or receiving an email stating that a user had saved some clips. If the email's receivers didn't have a BufferWare account, the message asked them to create an account and become a registered user. We designed BufferWare to be flexible in terms of use by individuals or groups. Interestingly, nearly every reported use of the system for archiving involved multiple people. People used BufferWare to share information. For example, users sent snippets of meetings when they wanted to save or share exact wording. For the most part, however, interviewees commented that such moments are rare and that current practices for documenting informal meetings, such as taking notes or photographing large shared writing surfaces, are often sufficient.

Almost all people who reported viewing archived video segments viewed more segments sent by other people than those they saved themselves. Even when they didn't expect an archive, people were intrigued when unknown video clips appeared in their accounts. Most unknown clips came from people who appeared to be experimenting with the system but chose not to register for an account.

Particularly in pervasive computing systems, where a technology's workings might be hidden, people often rely on those around them for indicators of a system's use and safety. So, it's important to exploit existing social practices when designing and deploying new pervasive computing technologies.

Experiential

People often use their past experiences with technologies to determine appropriate reactions to new technologies. Furthermore, a group's past experiences

that become shared cultural understanding impact responses to new technologies; people might already have beliefs about the "right" response within this culture. For example, previous experience with research projects contributed to people's respect for them and prevented even those individuals who felt negatively about recording from actively working against the BufferWare project. Past experiences with similar research projects also led many participants to believe that the researchers might examine the video clips, because past research projects had focused on such questions as the content of saved data.

The past experiences of this study's participants aren't representative of the general population. Many of the people impacted had computing-related backgrounds and interests. This group's understanding of computing adds new challenges, such as how to ensure that proper security mechanisms are in place for a group of people highly familiar with the potential security risks inherent in networked computing. It also created an environment in which people who are arguably knowledgeable about and con-

cerned with security and privacy issues could articulate these concerns.

Governmental and commercial entities have encountered other considerations when deploying systems in environments with technologically savvy users. These experiences provide a view into the potential range of reactions to new technologies, particularly as the world's population becomes more accustomed to computing and attuned to the issues

related to pervasive computing. Many solutions let concerned and technically inclined individuals manage and analyze risks for themselves. For example, the Linux open source community lets people compile their own operating systems. Similarly, when the wireless FastTrak system was installed to speed toll collection in Northern California's Bay Area, security and privacy details were made public so that those concerned could check the system details (see http://traffic.511.org/privacy.asp). Although this solution put some stakeholders at ease, in this and other pervasive computing environments in which a third party maintains service (such as toll collection or BufferWare), users must still trust that the inspected algorithms are the ones running.

Given these constraints, pervasive computing service providers must build the trust required for adoption and acceptance. Through repeated positive experiences with a particular service provider, users develop stronger feelings of trust, and the brand associated with that provider will likely become trusted. For example, the success of Google documents, spreadsheets, and mail demon-

A group's past experiences that become shared cultural understanding impact responses to new technologies; people might already have beliefs about the "right" response within this culture.

strates the Google brand's impact. People will store sometimes sensitive and personal information on the corporate servers largely because of the repeated positive experiences they've had with the search engine and other applications associated with the Google brand. In the BufferWare project, our research group's brand carried trust for certain stakeholders, enabling the system's adoption. However, rather than use a lengthy server

address and difficult-to-remember IP address, as is automatically assigned for university servers, we registered and used a short, easily remembered domain name. The divorce of the server address from a trusted set of server addresses associated with the Georgia Tech brand was a concern for some users.

Many stakeholders reported that recording in BufferWare was complicated and they didn't immediately understand it. They often ignored the notifications and descriptions provided and tried to reconcile this new mode of recording to their previous understandings. Rather than map BufferWare to the familiar and similar analog tape loops of security and other systems, they typically chose a different familiar technology—the streaming webcam. It isn't surprising, given this comparative technology, that they didn't immediately understand the automatic or even manual deletion of buffered or cached data. As designers, we must consider how much information people bring to their interactions with a captured space (for example, assumptions about use and retention policies), and find a straightforward way to educate and inform potential stakeholders, particularly when learning about a service isn't a high priority.

Not only is past experience with similar technologies relevant to understanding a new technology, but experiences with the new technology itself are also important. Giovanni Iachello and Jason Hong describe a "privacy hump" in which people start out largely pessimistic toward a new technology but become optimistic over time.[5] Repeated positive encounters with technology let people build experiential knowledge that leads to their accepting rather than fearing it.

At times, even after repeated exposure to a potentially risky technology has proved safe, people still have concerns, demonstrating an internal conflict between what they inherently believe and what they logically deduce. For example, one participant commented,

> There's the rational part of my brain that says … there's nothing nefarious going on. But then there's the paranoid part of my brain that says … wait, but they could be. How do I know they're not?

These conflicting sentiments indicate that there's often more to acceptability issues than risk-and-reward analyses or other rule-based decision criteria. Furthermore, people make decisions on a subconscious level that they can't articulate, creating a large design challenge.

A comparison of BufferWare's deployment with that of other pervasive computing technologies reinforces this conclusion. The deployment of a similar experimental "office memory" system at an Electricité de France research lab shows the effects of context.[6] EDF employees developed and used an audio-video recording system that continuously archived everything happening in the lab. Unlike BufferWare, the archiving was permanent, and all lab members could access the recordings. An interesting privacy control was that EDF tracked every access to the recordings, similarly to the optimistic security protocol,[7] and informed each individual of the identity of the person accessing the recordings of his or her workstation. These interactions' transparency let people build over time a set of experiences that led to an optimistic, safe model of the system's use.

S ervices in pervasive computing environments can perform actions for users without requiring much attention and effort. However, we must take care not to make "invisible computing" a literal and explicit goal. Instead, designers should leverage physical, social, and experiential knowledge to help users decide how to adopt and adapt to new pervasive computing technologies. As education about and experiences with technologies change over time, organizations should carefully plan the introduction of new technologies.[8] Understanding the three knowledge types can greatly enhance the design and adoption of pervasive computing technologies with privacy and security implications.

REFERENCES

1. D.A. Norman, *The Design of Everyday Things*, Doubleday, 1990.

2. G. Jancke et al., "Linking Public Spaces: Technical and Social Issues," *Proc. 2001 SIGCHI Conf. Human Factors in Computing Systems* (CHI 01), ACM Press, 2001, pp. 530–537.

3. D.H. Nguyen and E.D. Mynatt, *Privacy Mirrors: Understanding and Shaping Sociotechnical Ubiquitous Computing Systems*, tech. report GIT-GVU-02-16, Georgia Inst. of Technology, June 2002.

4. T. Erickson and W. Kellogg, "Social Translucence: An Approach to Designing Systems that Mesh with Social Processes," *Trans. Computer-Human Interaction*, vol. 7, no. 1, 2000, pp 59–83.

5. G. Iachello and J. Hong, "End-User Privacy in Human-Computer Interaction," *Foundations and Trends in Human-Computer Interaction*, vol. 1, no. 1, 2007, pp. 1–137.

6. S. Lahlou, "Living in a Goldfish Bowl: Lessons Learned about Privacy Issues in a Privacy-Challenged Environment," *Proc. Privacy Workshop at Ubicomp Conf.*, 2005, http://people.ischool.berkeley.edu/~jensg/Ubicomp2005/papers/2-Lahlou.pdf.

7. D. Povey, "Optimistic Security: A New Access Control Paradigm," *Proc. New Security Paradigms Workshop*, ACM Press, 1999, pp. 40–45.

8. G. Iachello et al., "Prototyping and Sampling Experience to Evaluate Ubiquitous Computing Privacy in the Real World," *Proc. 2006 SIGCHI Conf. Human Factors in Computing Systems* (CHI 06), ACM Press, 2006, pp. 1009–1018.

Gillian R. Hayes is an assistant professor in the Department of Informatics at the University of California, Irvine Donald Bren School of Information and Computer Sciences. Her research interests are human-computer interaction and ubiquitous computing, specifically in the areas of record-keeping and surveillance technologies. She received her PhD in computer science from the Georgia Institute of Technology. She's a member of the IEEE, ACM, and SigCHI. Contact her at the Donald Bren School of Information and Computer Sciences, Univ. of California, Irvine, 6210 Donald Bren Hall, Irvine, CA 92697-3425; gillianrh@ics.uci.edu; www.gillianhayes.com.

Erika Shehan Poole is a PhD student in the Georgia Institute of Technology's College of Computing and Graphics, Visualization, and Usability Center. Her research focuses on usability issues associated with networking and security. She received her BS in computer science from Purdue University. She's a member of the ACM and IEEE. Contact her at the GVU Center, Georgia Inst. of Technology, Technology Square Research Bldg., 85 5th Street NW, Atlanta, GA 30332-0760; erika@cc.gatech.edu; www.cc.gatech.edu/~erika.

Giovanni Iachello is an associate consultant at McKinsey & Co., focusing on high-tech and media strategy. His research interests relate to the human aspects of privacy and off-the-desktop technologies. He received his PhD in computer science from the Georgia Institute of Technology. He's a member of the IEEE, ACM, International Federation for Information Processing working group 9.6/11.7 "IT-Misuse and the Law," and SigCHI. Contact him at giac@clovermail.net.

Shwetak N. Patel is a PhD candidate in the Georgia Institute of Technology's College of Computing and Graphics, Visualization, and Usability Center. He's also an assistant director of the Aware Home Research Initiative. His research focuses on low-cost sensing systems, context-aware mobile phone applications, and technology to support ubiquitous computing applications. He received his BS in computer science from the Georgia Institute of Technology. He's a member of the IEEE Computer Society and the ACM. Contact him at the GVU Center, Georgia Inst. of Technology, Technology Square Research Bldg., 85 5th St. N.W., Atlanta, GA 30332-0760; shwetak@cc.gatech.edu; www.cc.gatech.edu/~shwetak.

Andrea Grimes is a human-centered-computing PhD student at the Georgia Institute of Technology. She studies the sociocultural nature of eating to design technologies that promote healthy eating in low-income urban communities. She received her BS in computer science from Northeastern University. She's a member of the ACM. Contact her at the GVU Center, Georgia Inst. of Technology, Technology Square Research Bldg., 85 5th St. NW, Atlanta, GA 30332-0760; agrimes@cc.gatech.edu; www.cc.gatech.edu/~agrimes.

Gregory D. Abowd is a professor in the Georgia Institute of Technology's School of Interactive Computing and Graphics, Visualization, and Usability Center, the director of the Aware Home Research Initiative, and the associate director of the Health Systems Institute. His research focuses on an application-driven approach to ubiquitous computing concerning both human-centered and technology-driven themes. He received his DPhil in computation from the University of Oxford. He's a member of the IEEE Computer Society, the ACM, and *IEEE Pervasive Computing*'s editorial board. Contact him at the GVU Center, Georgia Inst. of Technology, Technology Square Research Bldg., 85 5th St. NW, Atlanta, GA 30332-0760; abowd@cc.gatech.edu; www.gregoryabowd.com.

Khai N. Truong is an assistant professor in the University of Toronto's Department of Computer Science. His research lies at the intersection of human-computer interaction and ubiquitous computing, and focuses on usability and acceptance issues surrounding automated capture and access and context-awareness applications. He received his PhD in computer science from the Georgia Institute of Technology. Contact him at the Sanford Fleming Bldg., 10 King's College Rd. Rm. 3302, Toronto, ON M5S 3G4, Canada; khai@cs.toronto.edu; www.cs.toronto.edu/~khai.