An Ecosystem for Learning and Using Sensor-Driven IM Status Messages

The Nomatic prototype system and communications ecosystem automatically infers users' place, activity, and availability from sensors on their handheld devices or laptop computers and then reports this information to their instant-messaging contacts.

Donald J. Patterson, Xianghua Ding, Samuel J. Kaufman, Kah Liu, and Andrew Zaldivar
University of California, Irvine

The Internet is radically transforming itself in response to new perspectives that encourage community and group content creation while pushing technological dominance to the background. This Web 2.0 approach drives many high-traffic Web sites such as YouTube, Wikipedia, and the news-voting site Digg. Such sites utilize modest, easily duplicated technological innovations but are market leaders because of their large, active community of content creators and consumers.

In contrast, it's rarer to find sites that use passively generated content. Implicit voting (such as when people passively cast a vote for a video on YouTube by watching it) is the notable exception.

In our system, Nomatic (nomad and automatic),[1] we explore a way to amplify and leverage this passive style of content creation by focusing on status messages. These short bits of text are usually created by users in the context of small communities of people who monitor each other for play- or work-related distributed coordination. Status messages appear in instant-messaging (IM) clients as short customizable phrases such as "at lunch" or "out of the office." Commercial services also provide facilities for communicating status without IM (for example, Facebook, Twitter, and Jaiku).

By simply attaching sensor data to the status information that users enter in IM, we can create a rich ecosystem of context-aware applications that benefit the end user. At the most basic level, keeping status content up to date helps mitigate the increasing problem of interruptions in mobile communications, but there are many other potential uses of such data. To be effective, we must keep this ecosystem in balance by supporting the user's ability to provide status information, supporting other users' ability to understand that data, and effectively motivating both types of users to keep their status information accurate.

Structure and Sensors

Two unique ideas make our approach to status messages useful for passive content generation. First, we focus on mobile, structured, user-entered status updates. We're particularly interested in updates that people use to describe their activities when they're engaged in specific tasks while out using laptops and phones instead of desktop computers. Our user interfaces encourage update descriptions using fields labeled place, activity, and other. This approach is compelling because it aligns with mobile us-

Published by the IEEE CS ■ 1536-1268/09/$26.00

Figure 1. Nomatic uses sensors to provide instant-messaging (IM) contacts with user-generated contextual cues. These semantic interpretations are easier to keep up to date than custom status lines and easier to interpret than raw sensor data.

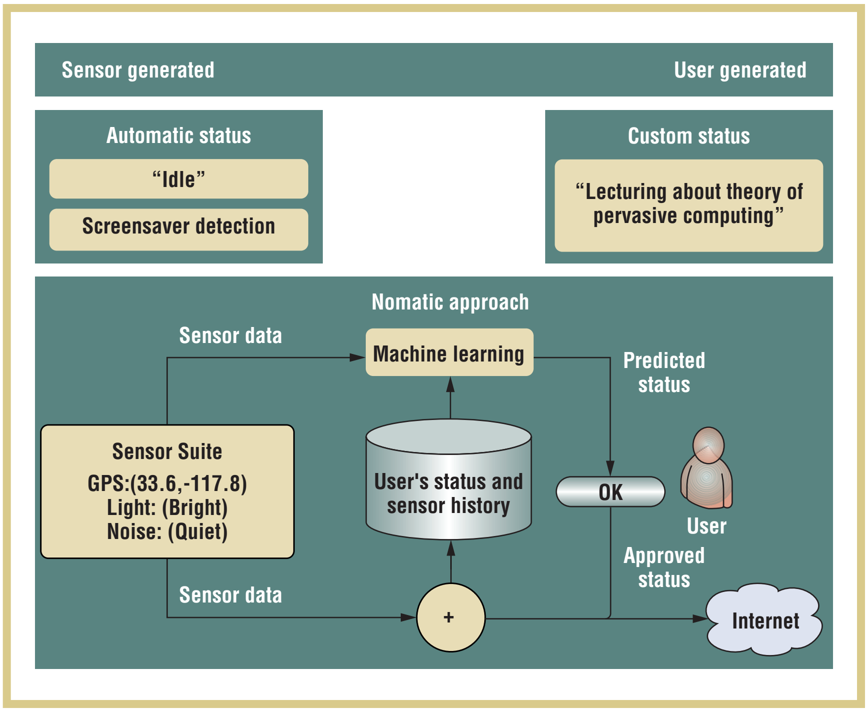

ers' existing practices,[2] reflects more directly on aspects of the physical world, and enables algorithms to make more assumptions about status content than would otherwise be appropriate.

Second, we rely on sensors built into commercial-off-the-shelf mobile platforms to help contextualize these status messages. By pairing structured status messages with sensor information such as Wi-Fi access points, ambient light levels, and accelerometer readings, it is possible to predict a user's status choice. Nomatic uses machine-learning algorithms to recognize situations in which sensor readings are similar to past situations and let users quickly re-enter those status messages later.

An Ecosystem

A viable passive content-generation system requires an ecosystem that supports the user's primary interest in status messages. To design such a system, we built on previous research that identified provision and perception as key aspects of an effective status ecosystem;[3] we then added motivation as a third factor. *Provision* draws attention to the fact that users must be able to easily and effectively describe their current status. *Perception* recognizes that other people must be able to view status for awareness to occur. *Motivation* emphasizes that for a user to keep status accurate, and for a computer to be able to interpret the status as content, it must be embedded in a needed and well-understood task.

Our theoretical orientation originates from Paul Dourish's idea of *embodied interaction*,[4] which emphasizes that pervasive computing is situated in a dynamic social and physical world that people are constantly renegotiating. We view status messages as digital probes into this negotiation. Rather than giving people a geo-tag tool and then making them guess about how to use such tags in the real world, our system observes users as they conduct their everyday activities and then uses that information as a building block for improving computational services.

Provision

Providing easy ways for users to enter their current status is important for making our status ecosystem viable. Most IM systems provide only two options. On one end of the spectrum, users must enter custom status messages. Custom status provides the most nuanced control over the presentation of a user's context, but keeping the messages up to date requires the user to repeatedly focus on the IM client. At the other end of the spectrum are status indicators that simply report raw sensor information that the audience is left to interpret. Typically, IM clients provide this service as an idle indicator based on a lack of keystrokes or the invocation of a screen saver. The first option produces intrusive but appropriate contextual information, while the second communicates accurate sensor cues that require interpretation. Our approach is to create a hybrid of these two extremes that supports rapid, easy updates and provides an interpretation of that information (see Figure 1).

Sensor Technology

Our machine-assisted approach uses sensors to provide rich descriptions of context through mobile status messages. This problem looks a lot like related research in place, activity, and availability detection, but by combining all three tasks in one application, we can provide a rich set of data for human observers to interpret. For example, knowing that users are in a classroom provides a great deal of information about their activities. Knowing that someone is teaching in a classroom provides even more information about whether he or she can be interrupted. To the best of our knowledge, no other researchers are investigating how to link these various aspects of context to help inform others.

Jeffrey Hightower and his colleagues

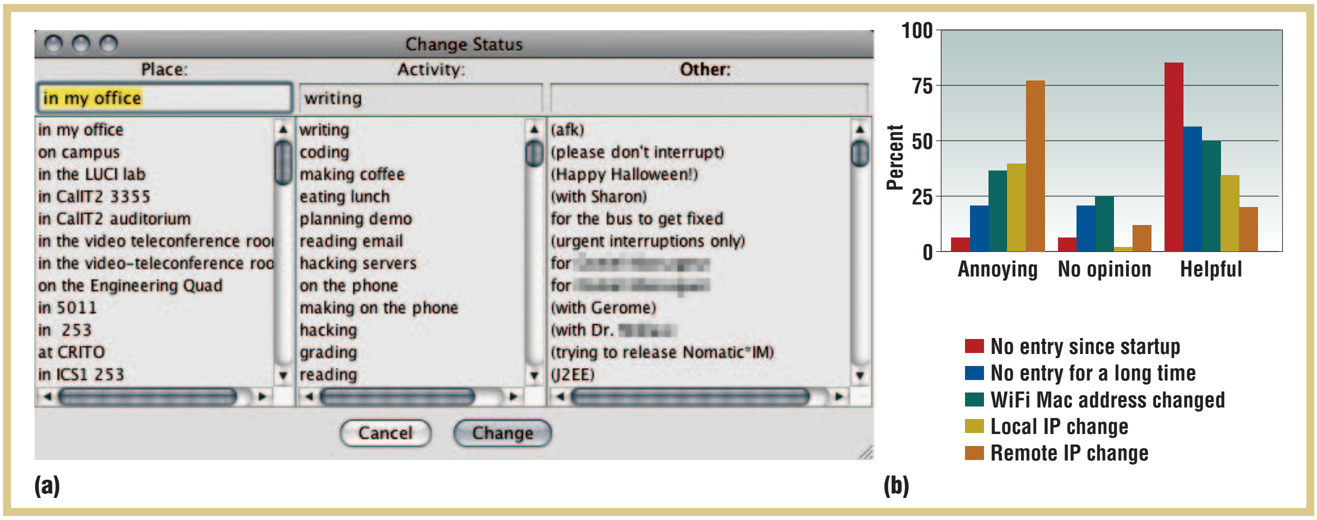

Figure 2. Nomatic in use. (a) A screenshot of a prototype machine-assisted status change interface; (b) percentage of time users rated proactive interruptions as helpful in response to various system events.

framed a portion of our problem as the *position-to-place problem*.[5] Their work focused on translating the exact, unambiguous sensor data from GPS streams (or Wi-Fi-based localization) into a more accurate description. This research brought several relevant challenges to light, particularly the many-to-many mapping of positions to places that make it difficult to determine correct place names. We address these concerns by attempting to determine an appropriate place name based on sensor data rather the "correct" place name. We validate our recommendations with the users' selections.

Our work is also closely related to *activity recognition*, in which many researchers are attempting to develop techniques to label a user's current activity from sensor streams. Much of this work attempts to characterize a user's sensor stream as being generated from 1 of *N* exclusive activities to support underlying reasoning algorithms.[6] We're unable to precategorize which activities users will be engaged in beforehand, so we take a different approach; we simply attempt to suggest the same activity to a user when they're in the same situation (as measured by the available sensors). This gives users control over how they wish to describe their activity, but limits the types of machine-learning approaches that are available to us.

Another closely related line of research is in detecting a user's availability from sensors.[7,8] Our desired application domain requires a more nuanced approach that gives users more flexibility than a numeric scale between interruptible and uninterruptible. For example, an on-call doctor in a restaurant might be interruptible for an emergency but not interruptible for a billing question. When the domain of interest moves away from the work environment, availability is much harder to define. Our approach is to push that decision into the social sphere. We attempt to learn a semantic description of the current sensor stream so that people can make a decision about whether it's appropriate to initiate a communication—that is, if callers knew that their callees were in a restaurant, they could self-censor themselves appropriately.

Obviously, we aren't the only ones working with context-enhanced IM systems. Other researchers are looking at how people respond to context-enhanced IM with regard to privacy (IMBuddy[9]). We're less concerned with privacy because we don't automatically broadcast sensor readings; we only broadcast status information that is predicted from sensor readings after the user has accepted them. Furthermore, our status information mimics the same language that the user used in a similar situation in the past, so users maintain control over their digital presentation. For example, if a user is in a coffee shop working and reports being "at work," we repeat that behavior even though an external location ontology would be unlikely to label Starbucks as "work."

The Awarenex system uses sensor modules to help support fluid conversation openings and closings,[10] which is similar to James Fogarty's work on interruption management.[11] We differentiate ourselves from these researchers by focusing on accurately predicting user-generated labels from the sensors built into current hardware. These labels might serve similar functions in conversations, but their function is less of a focus in our research.

Prototype Design

An effective user-friendly method of entering status requires that it be fast and easy for users to enter their status. We prototyped such an interface using a combination of machine-learning techniques and user interface design. The program flow works as follows (see Figure 1):

TABLE 1
The probability of various machine-learning algorithms correctly predicting place p, activity a, and other o status based on combinations of status and sensor ($\hat{s}$) readings.*

Algorithm	$P(p \mid \hat{s})$ (%)	$P(a \mid \hat{s})$ (%)	$P(a \mid p)$ (%)	$P(a \mid p, \hat{s})$ (%)	$P(o \mid \hat{s})$ (%)	$P(o \mid p, a)$ (%)	$P(o \mid p, a, \hat{s})$ (%)
Most likely	49	32	32	32	71	71	71
Decision stump	64	38	35	38	74	73	75
K nearest neighbors	75	50	**43**	53	74	77	77
Naive Bayes	85	72	**44**	73	78	**79**	82
Decision tree	85	63	**46**	64	81	**81**	84
Boosted stumps	**91**	76	**44**	80	85	**80**	91
Support-vector machine	**94**	**93**	**45**	**94**	**96**	**81**	**97**

* Bold numbers indicate best-performing combinations in each column that are statistically indistinguishable ($p < 0.05$).

- *Steady state.* While the user is working, our stand-alone program monitors the sensors on the computer and displays the current status in a small window.
- *Change status.* When the user presses a button to change his or her status, the interface reads the sensors on the computer and uses a machine-learning algorithm to make predictions for an appropriate status.
- *Validate predictions.* The interface displays the predictions in a selection box (see Figure 2a). Based on earlier results that indicated that 71 percent of all custom status messages in mobile IM are used to describe, place, activity, or availability,[2] we offer the user a template to enter place, activity, or a free-form label called "other". In the best case, the algorithm is correct, and users only have to accept the guesses by clicking on the "Accept Change" button. Otherwise, they can pick from the ranked list of guesses or, ultimately, type in a new status setting.
- *Accept prediction.* When the user accepts the new status, the interface again sweeps the sensors on the computer, pairs the sensor readings with the current status, and updates a local database that the machine-learning algorithm uses for future training.

Each of these steps provides ample opportunity for new lines of research, including understanding how to initiate a prompt when sensors change, choosing which algorithms are the best predictors, and experimenting with alternative user-interface designs.

Evaluation Methodology

To evaluate this interface's machine-learning aspect, we asked 14 mobile laptop IM users to try the status interface for three months. We then followed up during and after the study with short interviews to get feedback. We collected 7,154 status lines paired with sensor data that yielded an average of 5.7 status changes a day per person. From this information, we evaluated several classification strategies (see Table 1). We used two baseline classifiers: a static majority guess and a decision stump. Additionally, we tested three low-complexity algorithms: a K-nearest-neighbors algorithm ($K = 3$), a naive Bayes classifier, and a decision tree classifier. Finally, we tested two moderately complex algorithms: a boosted naive Bayes classifier and a support-vector machine (SVM). For each algorithm, we ran five runs of 10-fold cross-validation testing and training, using only a within-subjects prediction strategy.

For each of these machine-learning methods, we tested seven different prediction tasks. The features consisted of available sensors, $\hat{s}$ = { day of the week, time of day, local and remote IP address, Wi-Fi access point (AP) MAC address and service set identifier (SSID), currently active process, the number of displays connected to the computer and their resolutions, whether the mobile device was plugged in, 3D accelerometer readings, ambient light, computer's volume setting, and number of mouse clicks per second }. Depending on the task, we used place p, activity a, and other o as features or classification targets. From these features, we calculated $P(p \mid \hat{s})$, $P(a \mid \hat{s})$, $P(a \mid p)$, $P(a \mid p, \hat{s})$, $P(o \mid \hat{s})$, $P(o \mid p, a)$, and $P(o \mid p, a, \hat{s})$.

Analysis

By analyzing the results on our data set, we see that users have strong patterns of repeated status setting behavior in mobile IM (see Table 1). Based on the results from the single most likely classifier for place, $P(p \mid \hat{s})$, we can see that 49 percent of the time, users report being in their most frequent location, 32 percent of the time, they're doing their most frequent activity, and 71 percent of the time, they're using a single unique other status. Although these numbers are high from a temporal perspective, they mask various ways that people describe their places; in this case, the average number of unique places each user reported was 16. The most common places were custom variations on "at home" and "at work" but included descriptions as varied as "at the beach" and "in my car." The most common choice for activity status was to leave it blank, but that was closely followed by variations on activities such as "writing," "sleeping," and "working" and included "doing dishes" and "eating a chocolate muffin." Other o status was most often blank and after that was hard to categorize, with examples such as "George Lopez!!" "Don't bother me," "sick," and "Life's good."

Across the board, SVMs offered the best overall classification accuracy: they predicted the same place name that the user picked in 94 percent of the cases based on the sensors available, $P(p \mid \hat{s})$. We found similar results when predicting a user's activity, $P(a \mid \hat{s})$. The large gains made by the decision-stump algorithm over the most-likely algorithm in predicting place, $P(p \mid \hat{s})$, confirms that a few sensors contain a lot of information about place (such as Wi-Fi AP), but algorithms that leverage a variety of nonlocation-specific sensors do the best job of picking a nuanced place name. If we postulated knowing how people described their places, SVMs had a 45 percent chance of correctly guessing the associated activity, $P(a \mid p)$, which is a 13 percent improvement over just picking the most-likely activity as a default. Adding the user's choice of place names on top of the sensors available $P(a \mid \hat{s}) \rightarrow P(a \mid p, \hat{s})$ had negligible effect on improving activity prediction (93 percent → 94 percent). Finally, in predicting the other o status field, having knowledge of place and activity allowed SVMs to achieve 81 percent accuracy, $P(o \mid p, a)$, but the sensor information by itself was sufficient to improve that to 96 percent accuracy. As was the case with predicting activity, the additional interpretation by the user of the current place and activity $P(o \mid \hat{s}) \rightarrow P(o \mid p, a, \hat{s})$ had a negligible benefit (96 percent → 97 percent).

Because keeping status accurate is also a matter of keeping it up to date, we investigated strategies for interrupting users at moments when it appeared that their status might be changing. Our goal was to identify the most helpful times to interrupt a user. We programmed Nomatic to observe the sensor stream and look for various changes and, when they occur, to pop up the status change interface. After asking users to set their status, we also asked them to rate how helpful the reminder was. Figure 2b shows the interface, triggers, and results of the strategies.

From these data, we can see that users generally find it helpful to be prompted to change their status simply based on elapsed time, whether that occurs after start-up or after the user has been using the computer for a while. One effect that generated this result is that when a user takes a laptop out of standby, our interface is triggered. Intuitively, this seems to be a good time to get input from the user, and our data confirms this.

In contrast, interrupting the user during network change events isn't as clearly helpful. Although this seems to be a good indicator of a change in context related to mobility, our users experienced many situations in which their Wi-Fi connection rapidly switched between several competing APs while they weren't changing status. Similarly, intermittent Internet network problems made users' local and remote IP address change without a corresponding change in location, resulting in high annoyance scores for those triggers.

Perception

Using a status line to communicate aspects of your context isn't helpful unless the ecosystem also provides for perception of status by others, so we enabled Nomatic to interface with several third-party status broadcast services. Users can optionally choose to have their status reported to local IM clients (such as pidgin, iChat, and Adium) or popular microblogging services (such as Twitter and Facebook). Each of these services has mechanisms for broadcasting status to a social network of various types so that status can be used to negotiate appropriate responses to context.

Thinking about perception, we decided that our tool wouldn't automatically change status for a user. Because status communicates potentially sensitive information, we made this design decision because it's important for a user to trust our tool not to broadcast an inappropriate status. Nomatic requires that the user stay in the loop and click at least two buttons to initiate and then accept a status change. This exposes the relationship between accuracy and user attention that Nomatic is trying to minimize, but we hypothesize that remote viewers will find context status more useful if they can be confident that the user confirmed it. If contacts begin to doubt that a status is accurate, they're likely to ignore it, rendering it useless.

Another way of providing perception of status is through communal awareness displays. We designed such a display called Nomatic*Viz,[12] which visualizes the information that members of a small community can provide through our status interface. Because our software uses hardware that comes with existing computing platforms, rather than custom infrastructures, we were forced to address new ways of representing location in semantically meaningful ways when our users move outside of research environments. Instead of using geographic maps, Nomatic*Viz depicts historical and aggregate traces of participants' whereabouts in an abstract and ambiguous manner to convey the more useful context information that status provides, while attempting to stay away from the connotations of "tracking" or "monitoring" that raise privacy concerns.

Figure 3 shows a snapshot of Nomatic*Viz with 30 days' worth of data. In this visualization, the layout is dynamically determined by users' collective interactions with the Wi-Fi infrastructure. It depicts historical traces of people's whereabouts, but instead of using explicit icons for people, it uses different colors to distinguish between individuals. Therefore, participants and the people who engage with them

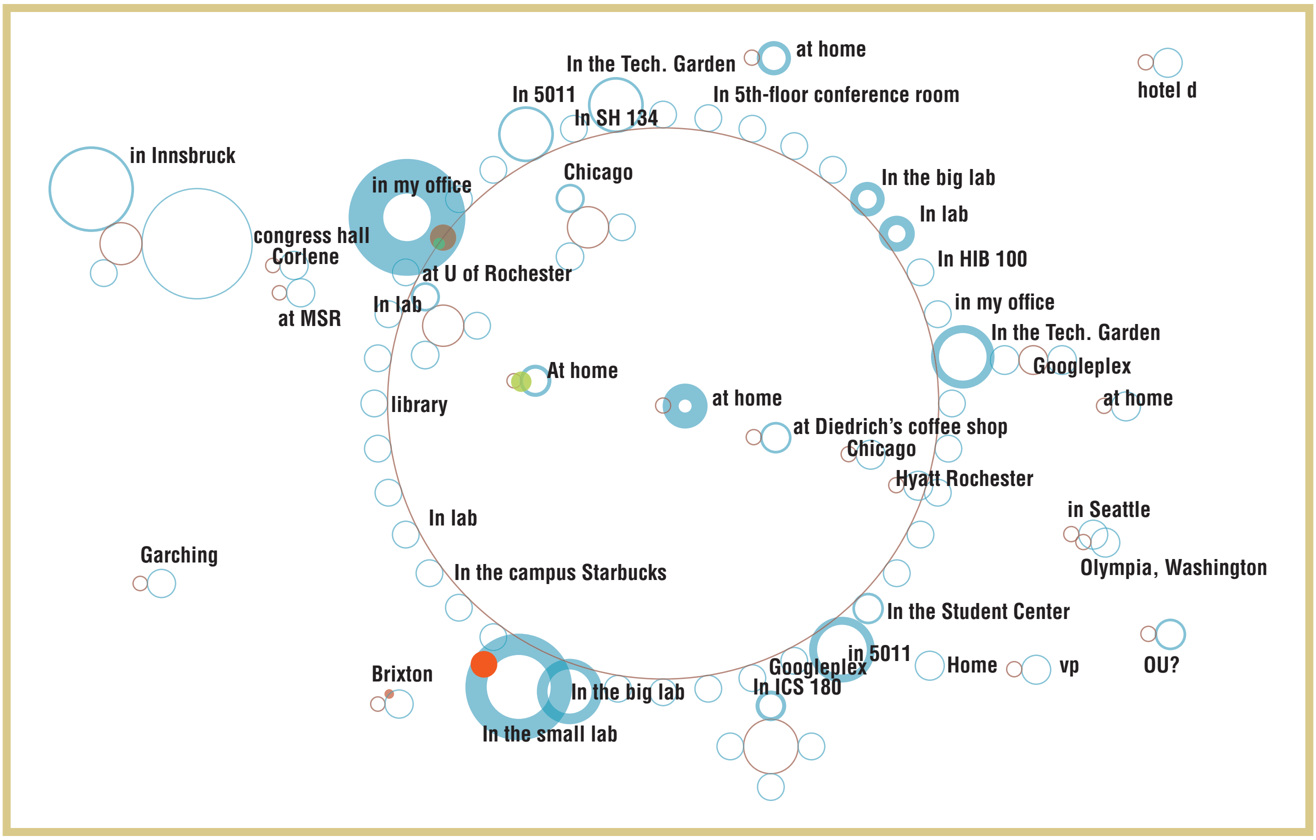

Figure 3. Nomatic*Viz displaying 30 days worth of data. The largest brown ring at the center is the University of California, Irvine, campus-wide SSID. The blue rings are Wi-Fi access points labeled with place names using our status interface. Blue ring size indicates the number of people who have visited the location, while border thickness indicates how often (in days). The brown ring clusters the Wi-Fi access points that have the same SSID; its position is determined by its size (the number of Wi-Fi hotspots associated) and the recency of the last visit. Colored dots indicate the most recent places from which users have reported.

and the display gradually learn how to interpret the information. Casual observers have limited insight into the details of a particular person's context.

Preliminary feedback of this display, which has been operating in our building's elevator lobby for more than a month, has been mixed. System users and viewers have revealed both an intuitive understanding of the abstract representations and confusion. On one hand, users' ability to interpret the visualization is correlated with their engagement with the community being visualized, which is a success for our design goals. Although viewers can't interpret all the traces, they can gain general impressions of the community's overall activities and easily recognize and interpret some of the participants' data.

On the other hand, the deployment's physical setting limits its effectiveness. Somewhat counter-intuitively, the elevator lobby is too transitional a space and wait times are too short for people to be able to develop a familiarity with what's going on in the visualization.

Motivation

Maintaining and monitoring status requires user effort. The ability to simply log status motivates some people to keep their status accurate—such recordkeeping is useful for managing billing records and for encouraging self-reflection. Additional motivating applications help further increase the benefits without additional effort, and subsequently improve the ecosystem for passive content creation.

IM Interruption and Embarrassment

Simply adding the ability for others to see your status creates *awareness*, a combination of provision and perception that researchers have thoroughly documented as having intrinsic value to support and improve distributed group work.[13,14] Studies have indicated that 13 percent of all pre-mobile IM dialog was simply related to negotiating availability.[15] As individuals are increasingly always online, IM has moved onto mobile platforms, and interruptions appear to be getting worse. More recent studies of mobile IM users have reported that 43 percent of them actively use strategies to manage interruptions and that as many as 7 percent have stopped using IM at one point or another due to the

the AUTHORS

Donald J. Patterson is the director of the Laboratory for Ubiquitous Computing and Interaction at the University of California, Irvine. His research interests are in the intersection of artificial intelligence, human-computer interaction, and ubiquitous computing. Patterson has a PhD in computer science from the University of Washington, Seattle. He is a member of the IEEE, ACM, and AAAI. Contact him at djp3@ics.uci.edu.

Xianghua Ding is a PhD candidate at the University of California, Irvine. Her research interests include human-computer interaction; computer-supported corporative work; ubiquitous computing; and designing, building, and studying interactive and collaborative applications in real settings. Contact her at dingx@ics.uci.edu.

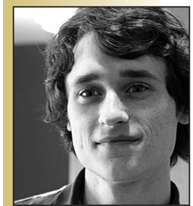

Samuel J. Kaufman is an undergraduate at the University of California, Irvine, majoring in informatics. His research interests include machine learning, human-computer interaction, and computer-supported cooperative work, including computer-mediated communication. Contact him at kaufmans@uci.edu.

Kah Liu is a master's student at the University of California, Irvine, and a technical director at Halcyon Design. His research interests include ubiquitous computing and human-computer interaction. Liu has a BS in information and computer science from the University of California, Irvine. Contact him at kahliu@gmail.com.

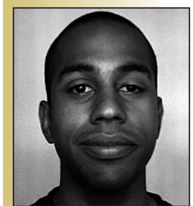

Andrew Zaldivar is an undergraduate in the Department of Cognitive Sciences at the University of California, Irvine, majoring in psychology. He is a member of the Social Code Group and a research assistant in the Cognitive Anteater Robotics Laboratory. His research interests include human-computer interaction, sustainability, and neuromodulation. Contact him at azaldiva@uci.edu.

distractions that it causes. Eighty percent of undergraduate mobile IM users report receiving embarrassing IMs because of their laptop screens' semipublic visibility.[2] We hypothesize that if machine-assisted status messages effectively mitigate interruptions and alert users about the social context of a contact before they initiate, then it's a significant motivation for keeping status up to date.

CATDL

We can also incorporate status information into a context-aware to-do list (CATDL). By reordering a long list of tasks to focus attention on those that are currently relevant via a process that the user is already engaged in (status setting), we make context maintenance even more valuable. This approach enables a grocery list to appear on a mobile device upon entering a grocery store and a list of work tasks to naturally shift to a list of home tasks at the end of a work day.

We've implemented such a system as a locally hosted Web application called Nomatic*CATDL, which combines the status from Nomatic with an online to-do list manager called Remember the Milk (www.rememberthemilk.com). When users want to see their current to-do list, they open the Web application, which establishes a local connection to the Nomatic status setting tool and a remote connection to Remember the Milk. Nomatic contextually sorts the task list in the browser and presents it as a lightweight user interface suitable for a mobile phone.

Geographic Overlays

Although the Nomatic*Viz display is highly abstract, there are also times when it's appropriate to more explicitly map a context's geographic location. In public or professional spaces, people often explicitly reveal or are required to reveal their current task. In such situations, it might be effective to explicitly represent status as a function of position. We're exploring this idea through the design and study of a touchscreen directory kiosk for a large academic building. Existing kiosks help people find location according to facility managers' descriptions (such as building names and room numbers) but don't reveal colloquial place names, such as a familiar class names (for example, Physics for Poets).

We've prototyped a display that overlays a traditional building map with historical and real-time status information (see Figure 4). In this design, users can select a temporal range of data that they want to visualize at the top. This selection populates three drop-down lists on the lower left with data that other users have previously entered through the Nomatic tool. This building has been calibrated to support localization by Wi-Fi fingerprinting, so status can be associated with a particular geographic location. To support anonymity, Nomatic only displays the status entries duplicated by more than one person. The display's user can then select a place, activity, or other status from the drop-down list; it highlights locations at which that status label has already been used with colored circles on the map. We're in the process of studying how the kiosk's enhanced design affects user interaction with the directory display.

Figure 4. A context-enhanced directory for finding events in real time. An overlay displays a user-selected temporal range of data.

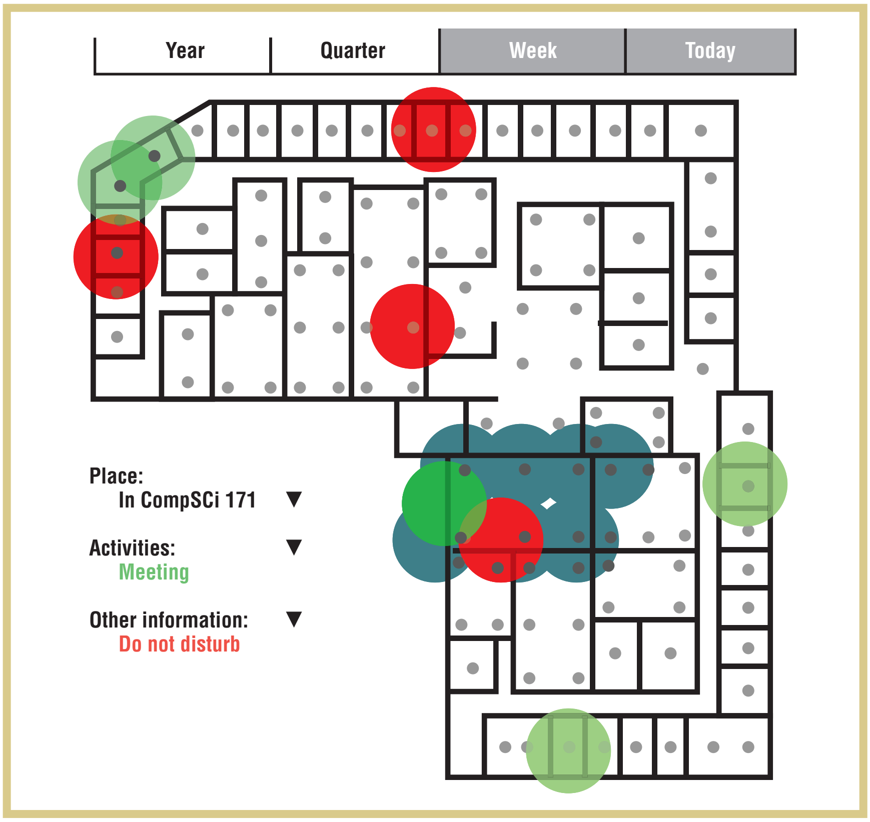

As we continue to explore possibilities and identify how to innovate existing technologies to best support users' needs, we see many opportunities for future research. Interesting work remains to be done in developing a more sophisticated understanding of when users would like to be interrupted to change their status. There are also opportunities for developing mass-collaboration approaches for sharing sensor and status information across users. Finally, we could identify which sensors are the most valuable for setting status as a way of advising future hardware construction. This includes creating abstract virtual sensors based on users' calendars, their contacts statuses, or any number of digital probes that might influence their mood.

REFERENCES

1. D.J. Patterson, X. Ding, and N. Noack, "Nomatic: Location by, for, and of Crowds," M. Hazas, J. Krumm, and T. Strang, eds., *Proc. 2nd Int'l Location- and Context-Awareness Workshop* (LoCA 2006), LNCS 3987, Springer, 2006, pp. 186–203.
2. D.J. Patterson et al., "Online Everywhere: Evolving Mobile Instant Messaging Practices," J. McCarthy, J. Scott, and W. Woo, eds., *Proc. 10th Int'l Conf. Ubiquitous Computing* (UbiComp 2008), ACM Press, 2008, pp. 64–73.
3. G. Fitzpatrick, *Locales Framework: Understanding and Designing for Wicked Problems*, Kluwer Academic, 2003.
4. P. Dourish, *Where the Action Is: The Foundations of Embodied Interaction*, MIT Press, 2001.
5. J. Hightower, "From Position to Place," *Proc. 2003 Workshop on Location-Aware Computing*, 2003, pp. 10–12; http://research.microsoft.com/en-us/um/people/jckrumm/publications%202003/2003workshoponlocationawarecomputing.pdf.
6. D.J. Patterson et al., "Fine-Grained Activity Recognition by Aggregating Abstract Object Usage," K. Mase and B. Rhodes, eds., *Proc. IEEE Int'l Symp. Wearable Computers* (ISWC), IEEE CS Press, 2005, pp. 44–51.
7. J.B. Begole, N.E. Matsakis, and J.C. Tang, "Lilsys: Sensing Unavailability," *Proc. ACM Conf. Computer-Supported Cooperative Work* (CSCW 04), ACM Press, 2004, pp. 511–514.
8. E. Horvitz, P. Koch, and J. Apacible, "Busybody: Creating and Fielding Personalized Models of the Cost of Interruption," *Proc. ACM Conf. Computer-Supported Cooperative Work* (CSCW 04), J.D. Herbsleb and G.M. Olson, eds., ACM Press, 2004, pp. 507–510.
9. G. Hsieh et al., "Field Deployment of IMBuddy: A study of Privacy Control and Feedback Mechanisms for Contextual IM," J. McCarthy, J. Scott, and W. Woo, eds., *Proc. 10th Int'l Conf. Ubiquitous Computing* (UbiComp 08), ACM Press, 2008, pp. 91–108.
10. J.C. Tang, "Approaching and Leave-Taking: Negotiating Contact in Computer-Mediated Communication," *ACM Trans. Computer-Human Interaction*, vol. 14, no. 1, 2007, p. 5.
11. J. Fogarty, J. Lai, and J. Christensen, "Presence versus Availability: The Design and Evaluation of a Context-Aware Communication Client," *Int'l J. Human-Computer Studies*, vol. 61, no. 3, 2004, pp. 299–317.
13. C. Heath and P. Luff, "Disembodied Conduct: Communication through Video in a Multimedia Office Environment," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 91), ACM Press, 1991, pp. 99–103.
14. J.C. Tang et al., "Connexus to Awarenex: Extending Awareness to Mobile Users," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 01), ACM Press, 2001, pp. 221–228.
15. M. Handel and J.D. Herbsleb, "What Is Chat Doing in the Workplace?" *Proc. 2002 ACM Conf. Computer-Supported Cooperative Work* (CSCW 02), ACM Press, 2002, pp. 1–10.
12. X. Ding and D.J. Patterson, "Status on Display: A Field Trial of Nomatic*Viz," *Proc. European Conf. Computer-Supported Cooperative Work* (ECSCW 09), I. Wagner et al., eds., Springer, 2009, pp. 303–322.