



Nigel Davies  
Lancaster University

Daniel P. Siewiorek  
Carnegie Mellon University

Rahul Sukthankar  
Intel Research

**A**ctivity-based computing has its roots in research into context-aware systems. Such systems take into account the user's state and surroundings and enable the user's mobile computer or environment to adapt accordingly. The wearable and pervasive computing communities, among others, have explored context-aware computing in detail as part of the drive to simplify HCI.

### Minimizing user interaction

Wearable or body-worn computers provide hands-free operation and offer compelling advantages in many applications. Wearable computers deal with information rather than programs, becoming information tools in the user's environment—much like a pencil or a reference book. They provide portable access to information and can automatically accumulate information as the user interacts with and modifies the environment. This eliminates the costly and error-prone process of transferring information to a central

database. At the core of these ideas is the notion that wearable computers should seek to merge the user's information space with his or her workspace, blending seamlessly into the existing work environment to provide as little distraction as possible.

While wearable computing research focuses on developing mobile devices that users can carry, pervasive computing focuses on embedding computers into the environment. Pervasive computers aim to merge into the fabric of everyday life to become almost invisible. Readers of this magazine will be familiar with pervasive computing and the many forms that it can take. (If not, see Mark Weiser's seminal article, "The Computer for the 21st Century," reprinted in this magazine's premier issue.)

The unifying factor for both wearable and pervasive computing systems is a desire to minimize explicit user interaction—to either reduce distractions or help the computers blend into the background. Context-aware computing provides mechanisms for achieving this.

### Context-aware computing

Initially, context-aware systems used location as the principal form of context. However, adding low-cost sensors (such as for measuring acceleration, light, and audio) to mobile and pervasive platforms, combined with advances in machine learning, enables systems to build a much richer model of the user's context.

For example, body-mounted accelerometers can recognize a variety of human activities—from common activities such as walking or sitting to higher-level activities such as driving a car or riding a bus. They can even distinguish between arm motions for different manual-wheelchair propulsion styles and determine the surface on which the wheelchair is traveling (such as a rug, tile, or asphalt). A user's context can thus be quite rich, consisting of attributes such as physical location, physiological state (such as body temperature, heart rate, and skin resistance), emotional state, personal history, and daily behavioral patterns.

The articles in this special issue fo-

cus on context-aware systems that recognize activities. Such systems interpret sensor data as a reflection of users' behaviors and actions as part of an activity. Understanding the activity involved lets systems provide computational support at a familiar level of abstraction—namely, in terms of tasks and activities rather than in terms of low-level sensor data or events. For example, activity-based systems might be able to identify that a user is “chopping vegetables while cooking a meal” rather than simply identifying that the user is in the kitchen, moving his or her hands in a specific manner. This activity knowledge can help systems adapt their services and computational resources to a user's context, improve a user's ability to collaborate with others, and help the user stay focused on the task at hand.

Numerous potential applications exist for activity recognition systems in domains as diverse as assisted living, remote worker support, and homeland security. Many research projects have explored the topic, including attempts to create instrumented environments and wearable sensors that can capture low-level data about user actions. Researchers have also studied various techniques for recognizing activities using statistical methods and have explored how users can specify activities and interact with activity-aware systems.

### In this issue

The first article in this issue, “Rapid Prototyping of Activity Recognition Applications,” describes the Context Recognition Network (CRN) Toolbox for rapid prototyping of context and activity recognition applications. Parameterizable and reusable software components provide an array of algorithms for multimodal sensors, signal processing, and pattern recognition. The CRN Toolbox should significantly reduce the development time associated with ac-



**Nigel Davies** is a professor of computer science at Lancaster University and an adjunct associate professor of computer science at the University of Arizona. His research interests include systems support for mobile and pervasive computing. He focuses on the challenges of creating deployable mobile and ubiquitous computing systems that can be used and evaluated “in the wild.” He's an associate editor of *IEEE Pervasive Computing*. Contact him at [nigel@comp.lancs.ac.uk](mailto:nigel@comp.lancs.ac.uk).



**Daniel P. Siewiorek** is the Buhl University Professor of computer science and electrical and computer engineering at Carnegie Mellon University. He's also director of the Human-Computer Interaction Institute. His research interests include systems architecture, reliability, modular design, wearable computers, and context-aware computing. He received his PhD in electrical engineering from Stanford University. He's a member of the IEEE Computer Society, the ACM, Tau Beta Pi, Eta Kappa Nu, and Sigma Xi. Contact him at [dps@cs.cmu.edu](mailto:dps@cs.cmu.edu).



**Rahul Sukthankar** is a senior principal research scientist at Intel Research and an adjunct research professor in robotics at Carnegie Mellon University. His current research focuses on computer vision and machine learning, particularly in the areas of object recognition and information retrieval in medical imaging. He received his PhD in robotics from Carnegie Mellon. Contact him at [rahuls@cs.cmu.edu](mailto:rahuls@cs.cmu.edu).

tivity recognition systems, thus facilitating the emergence of many more prototype activity-based applications.

The next article, “The Mobile Sensing Platform: An Embedded Activity Recognition System,” describes the evolution and design of a four-ounce wearable mobile sensing platform that incorporates multimodal sensing, data processing, machine learning, and wireless communications for activity recognition.

While the first two articles focus on underlying technologies for developing activity-aware systems, the remaining two articles describe experiences with activity recognition in real-world applications. “Wearable Activity Tracking in Car Manufacturing” explores multisensor systems for training new workers on complex assembly tasks and for notifying quality-control inspectors when they've missed an inspection task. It details several practical issues related to deploying new technology in a production environment and will be of interest to anyone considering deploying an activity-tracking system in

an industrial setting.

Finally, “Activity-Aware Computing for Healthcare” discusses design principles for developing robust activity recognition in hospital applications. The authors have deployed a prototype system in a clinical setting and show how activity-aware computing could significantly improve the efficiency of medical professionals in healthcare.

Despite the wealth of work in this area, activity recognition remains a fertile research area with many unanswered questions. How should activities be specified, and who should do this? How many activities can systems robustly detect, and how should they inform users of their decision-making progress? How can we harness activity-aware systems to help in diverse application domains such as healthcare and remote working? We hope that this special issue helps advance the state of the art in this key area of mobile and pervasive computing research. ■