

# Opportunistic Human Activity and Context Recognition

**Daniel Roggen and Gerhard Tröster**, *ETH Zurich, Switzerland*

**Paul Lukowicz**, *DFKI and Kaiserslautern University of Technology, Germany*

**Alois Ferscha**, *Johannes Kepler University of Linz, Austria*

**José del R. Millán and Ricardo Chavarriaga**, *École Polytechnique Fédérale de Lausanne, Switzerland*

---

**Achieving true ambient intelligence calls for a new opportunistic activity recognition paradigm in which, instead of deploying information sources for a specific goal, the recognition methods themselves dynamically adapt to available sensor data.**

---

**M**ore than 20 years ago, Xerox PARC's Mark Weiser laid out a vision in which highly miniaturized wireless sensor nodes and computers, so small that they would disappear in the background, would assist us in all facets of our lives.<sup>1</sup> By being ubiquitous, they would know what we need at all times and deliver support proactively.

This vision of *ambient intelligence* continues to drive researchers. The Internet of Things (IoT) now provides the necessary infrastructure to transparently access sensors, processors, and actuators using standardized protocols regardless of hardware, operating systems, or location.<sup>2</sup> For example, the IPv6 protocol forms the IoT's backbone.

To realize ambient intelligence, these "things" must

understand the user's context,<sup>3</sup> including the location, activities (gestures, body posture, modes of locomotion), cognitive/affective states, and social interactions as well as the environment's state. Knowing this information makes it possible to provide unprecedented types of user support. Consider, for example, a person with dementia who forgets to pour water into a container on a hot plate in the kitchen. Activity recognition could help a pervasive monitoring system identify this situation and then either inform the user or turn off the hot plate. Many other high-risk or potentially dangerous scenarios at home, in the workplace, and other locations would likewise benefit from this capability.

However, current context-recognition approaches are restricted to domains with dedicated sensors. Achieving true ambient intelligence calls for a new *opportunistic activity recognition* paradigm. Instead of deploying information sources for a specific recognition goal, the methods themselves must adapt to the data available at any time. We present novel techniques that allow for opportunistic activity recognition in dynamic sensor configurations.

## A NEW PARADIGM

The IoT provides access to many sensors "for free" as

a by-product of other applications. On-body sensors are found in smartphones, wristwatches and other wearable or portable gadgets, as well as in smart clothing and shoes. Homes and offices include a rich array of sensors for lighting and climate control, appliance usage (oven and refrigerator door sensors), electricity usage metering, security (motion-activated lights and surveillance cameras, door/window movement sensors for intrusion detection), and entertainment (motion-sensing game consoles like Microsoft Kinect for Xbox 360 and the Nintendo Wii U). The number of sensors in our environments is likely to increase in the future. As users change locations, pick up or leave behind devices, change smart clothing, or interact with smart objects, they will encounter various sensor configurations in their surroundings.

Despite this potential for widespread sensing, however, ambient intelligence is still far from ubiquitous. Indeed, as the “Classic Activity Recognition Paradigm” sidebar explains, researchers derive the probabilistic models that current signal processing and machine learning techniques use to identify activities in sensor data streams from datasets collected at design time with predefined and optimal sensor configurations. However, collecting datasets for all possible sensor configurations is clearly not feasible. A new opportunistic activity recognition paradigm is thus required to realize ambient intelligence with readily available resources.

## BRIDGING THE ABSTRACTION GAP

Service-oriented architectures (SOAs) have become a common abstraction for building context-aware applications.<sup>4</sup> In the IOT, smart things advertise their services, such as recognizing a contextual element—for example, presence in a room. SOAs can include autonomic properties<sup>5</sup> and adapt to changing resources,<sup>6</sup> which makes it possible to substitute service providers and compose services to address complex problems. At the core of the approach lies a semantic description of application goals, service capabilities, and their interrelations, enabling semantic reasoning. Many context-aware middleware systems support SOAs in dynamic environments.<sup>7</sup>

SOAs require that sensor nodes provide contextual data. This is relatively straightforward with RFID, presence, or temperature sensors. An SOA can apply reasoning on single sensor readings and easily substitute semantically identical sensors, such as one temperature sensor for another in the same room. However, these sensors cannot understand complex manipulative gestures, body postures, or modes of locomotion.

Opportunistic activity recognition requires the ability to detect meaningful patterns hidden in noisy time series spread across multiple on-body and environmental sensors. Advanced signal processing and machine learning techniques are needed to recognize these patterns. In

## Classic Activity Recognition Paradigm

**A**n activity recognition system exploits the fact that sensors—whether ambient, object, or wearable—deliver characteristic raw signal patterns when a user carries out an activity or encounters a specific situation.

System designers first deploy sensors according to the activities to recognize—for example, opening a door, cooking, or eating. A common on-body sensor is the accelerometer, which is cheap, small, and delivers rich information. Typically placed on the wrist, it provides unique signals for gestures as diverse as drinking, turning a knob, or pulling a drawer. More complex activities require combinations of sensors.<sup>1</sup>

Designers then devise a recognition chain to detect the characteristic signal patterns. To support robustness and computational complexity, the system relies on signal processing and machine learning techniques.<sup>2,3</sup> It first applies preprocessing to remove noise from the signals and segment them. Next, the system extracts signal features to enhance the characteristics unique to each activity and to reduce data dimensionality, then it uses classifiers to map these features to discrete activity or context classes. If the system uses many sensors, it can fuse the decisions of several classifiers operating on individual sensors. The system often performs individual classification on the nodes to reduce network load, fusing decisions centrally. At this stage, the system calculates class occurrence probabilities to recognize more complex activities through probabilistic reasoning.

The system obtains recognition chain parameters—cutoff frequencies, kinds of features, classifier properties, and so on—by training on an annotated dataset of sensor readings recorded as users execute the activities of interest. Larger datasets capture greater variability in activity execution and sensor characteristics.

Although classic activity recognition uses probabilistic approaches, it is optimized to the selected sensors. Changing the sensors requires redesigning the system.

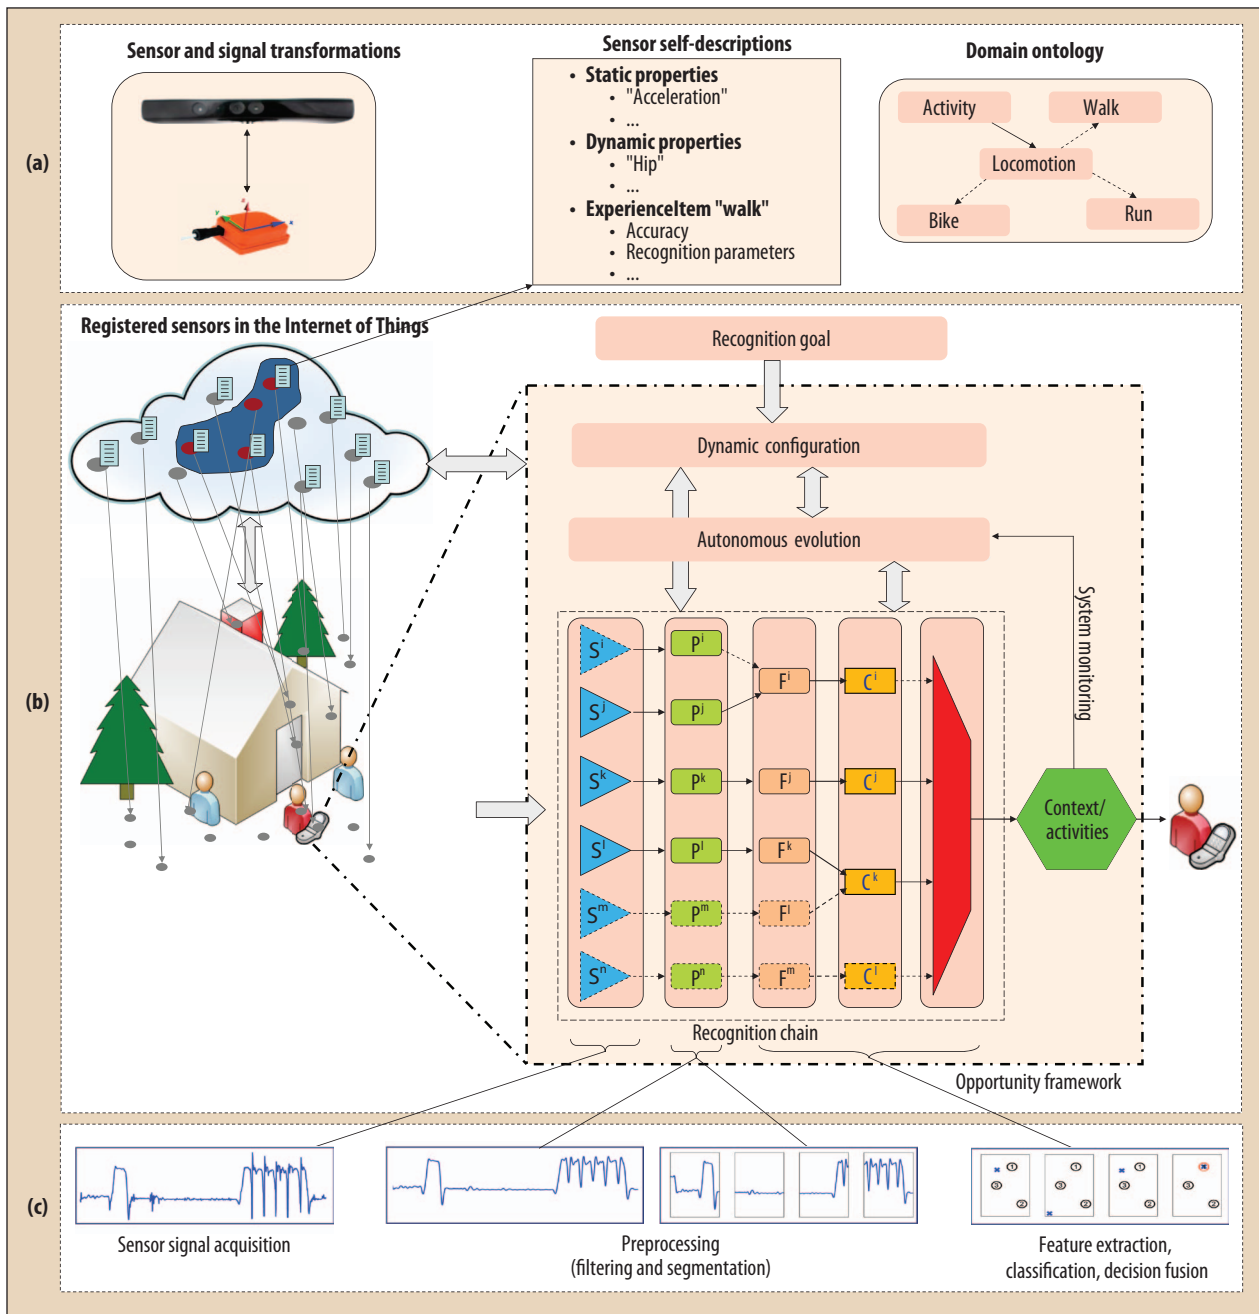
### References

1. T. Stiefmeier et al., “Wearable Activity Tracking in Car Manufacturing,” *IEEE Pervasive Computing*, vol. 7, no. 2, 2008, pp. 42-50.
2. S.J. Preece et al., “Activity Identification Using Body-Mounted Sensors—A Review of Classification Techniques,” *Physiological Measurement*, Apr. 2009, pp. R1-R33.
3. D. Figo et al., “Preprocessing Techniques for Context Recognition from Accelerometer Data,” *Pervasive and Mobile Computing*, vol. 14, no. 7, 2010, pp. 645-662.

addition, identical types of sensors cannot be substituted naively for one another, as a sensor’s location, orientation, and field of view influence its readings—for example, a smartphone’s motion sensor delivers very different signals when its orientation changes, or if it is placed in a shirt versus a trousers pocket. Opportunistic activity recognition must account for this highly variable input.

## OPPORTUNISTIC ACTIVITY RECOGNITION

We envision ambient intelligence as an autonomous



**Figure 1.** Opportunistic activity recognition system architecture. (a) The system leverages domain knowledge including sensor self-descriptions, signal and sensor transformation rules, and an activity ontology. (b) The system dynamically configures an activity recognition chain based on the recognition goal, the available sensors, and the domain knowledge. It continuously analyzes relations between sensors and activities to expand the domain knowledge, thus realizing autonomous evolution. (c) The recognition chain maps sensor signals ( $S$ ) to activity or context classes by applying preprocessing ( $P$ ), feature extraction ( $F$ ), classification ( $C$ ), and decision fusion methods. The solid arrows indicate streaming dataflows, while dashed arrows indicate sensors that might appear or disappear. The signal example is acceleration data measured at the wrist. The second half of the signal corresponds to the activity "cleaning table," characterized by periodic hand movements.

system, in the artificial intelligence (AI) sense, with the capacity to solve new problems, adapt online, analyze its own behavior, and learn from past experience. The novelty of our approach is that it adapts the data-processing steps alongside the sets of sensors used. A combination

of advanced streaming signal processing and machine learning techniques with sensor self-descriptions and a management framework make it possible to solve recognition problems in a wide range of situations using opportunistic sensing.

Only a few situations change the opportunistic configuration of sensor nodes:

- *Sensor appearance.* The user enters a sensor's range, resulting in the appearance of one or more signal channels. The sensor is *known* if it was expected at design time or has already been detected. It is *unknown* if it was unforeseen or has been detected for the first time and no metadata exist to immediately use it—this could be the case for legacy devices introduced in the IoT.
- *Sensor disappearance.* The user leaves a sensor's range or the sensor is powered down, resulting in the loss of signals.
- *Change in sensor characteristics.* Various conditions can change sensor signal patterns. For example, power-management adjustments to the signal sampling rate, displacement of an on-body sensor along a limb, or sensor degradation due to harsh conditions could cause a moderate change, while displacement of an on-body sensor to a completely different body location, an alteration of the field of view of ambient sensors, or sensor failure could produce a major change.

In these situations, an opportunistic activity recognition system must efficiently use the available resources and keep working when the sensor configuration changes; ideally, it should be capable of improving throughout these changes.

As Figure 1 shows, our system achieves this with a sensing and context framework that manages an adaptive *recognition chain*. The sidebar “The Opportunity Framework” describes the framework in detail. The recognition chain is made up of interconnected streaming signal processing and machine learning elements that map the raw sensor signals to activity primitives. To capitalize on the vast number of multimodal sensor nodes in the IoT, the recognition chain uses dynamic classifier ensembles in which each node has a dedicated pipeline for signal preprocessing, feature extraction, and classification. The opportunity activity recognition system then fuses the decisions of the node-specific classifiers into a global decision by, for example, majority voting. This architecture allows local processing on the nodes and is well suited to handle node addition and removal.

The system generates solutions to a recognition problem on the fly using domain knowledge, dynamic configuration, and autonomous evolution.

## Domain knowledge

Just as the availability of domain knowledge enables an AI system to exhibit more advanced behaviors, our opportunistic activity recognition system relies on rich domain

## The Opportunity Framework

**O**ppportunity is a reference framework for an opportunistic activity recognition system designed around a planner.<sup>1</sup> The system autonomously executes a recognition chain in the form of a directed dataflow graph of sensor data, relying primarily on signal processing and machine learning techniques to detect different activities or contexts.

The framework includes a set of descriptors and tools to generate recognition-chain configurations on the fly. This includes an activity-related ontology, self-descriptions in markup language derived from SensorML that the sensor nodes must deliver once they register with the system, and sensor and signal transformations.

Self-descriptions can be updated dynamically to reflect new node properties discovered by opportunistic methods. In a prototypical scenario involving the recognition of daily activities, we related about 200 low-level activities to 15 high-level activities. Self-descriptions advertise the activities that the sensors can recognize and the recognition-chain parameters. Upon execution of the recognition chain, the system attempts to match the recognition goal with the sensors' capabilities. It uses goal reasoning and sensor substitutions to find a configuration solving the recognition problem.

The framework is implemented using Java and OSGi modules and can be deployed on various platforms, including embedded ARM boards with capabilities identical to those of current smartphones. An OSGi module encapsulates each sensor abstraction and processing element. The framework supports centralized or decentralized execution of the recognition chain. The system accesses sensors in the Internet of Things over TCP/IP.


### Reference

1. M. Kurz et al., “The OPPORTUNITY Framework and Data Processing Ecosystem for Opportunistic Activity and Context Recognition,” *Int'l J. Sensors, Wireless Comm. and Control*, vol. 1, no. 2, 2012, pp. 102-125.

knowledge to identify solutions to recognition goals.

Sensor nodes store part of this knowledge as self-descriptions. The system caches this information upon node discovery, obtaining an overview of the capabilities of nodes in the user's surrounding. A sensor self-description contains static parameters, such as the physical quantity that the node measures. It also contains ExperienceItems, each of which indicates a goal toward which the node can contribute (for example, “walk”), when this is possible (for example, “body placement” is required for “walk”), the parameters of the recognition chain required to achieve that goal (for example—features, classifier), and the recognition accuracy.

We provide an initial set of ExperienceItems at design time, and train these in the same way as is done with the classic paradigm. For instance, “walk” could be derived from a gait dataset and indicate that the energy in a 1- to



**Figure 2.** The opportunistic activity recognition system leverages a wide range of sensors available in the user's surroundings. As the user moves around, he enters and leaves the sensing range of these devices. This scenario includes five typical situations: (1) coming home, (2) relaxing, (3) cooking, (4) gaming, and (5) going jogging.

3-Hz frequency band for a foot-mounted accelerometer is a feature suited to a decision tree classifier.

An ontology relates activity concepts to support goal reasoning, and known sensor transformations indicate how sensors can be combined to emulate others. We provide the system with basic transformations for typical modalities. For example, the system can process signals from a compass to emulate a gyroscope. This allows rich capabilities for substituting sensors.

### Dynamic configuration

As in SOAs, our approach favors dynamic configurations. We specify what the system is to recognize, but not how. At runtime, the system searches for a configuration that satisfies the recognition goal. It uses goal reasoning to restate the goal according to the nodes' capabilities. Because some required sensors might not be available, the system also attempts to match available and required sensors through sensor transformations.

### Autonomous evolution

As with AI learning agents, our system can adapt and expand its domain knowledge as new situations arise at runtime. This autonomous process combines self-monitoring, -adaptation, and -learning.

Nodes can self-monitor and describe dynamic parameters, such as their current orientation, and then update dynamic entries in their self-description. The system can autonomously adapt ExperienceItems or create new ones when it discovers new solutions to a recognition goal. It also can store new sensor transformations as it discovers them.

Over time, the system evolves to operate in unforeseen situations beyond the initially provided domain knowledge.

### SYSTEM OPERATION

Following a user as he moves about his house illustrates how the opportunistic activity recognition system operates. Figure 2 focuses on five typical situations: coming home, relaxing, gaming, and going jogging.

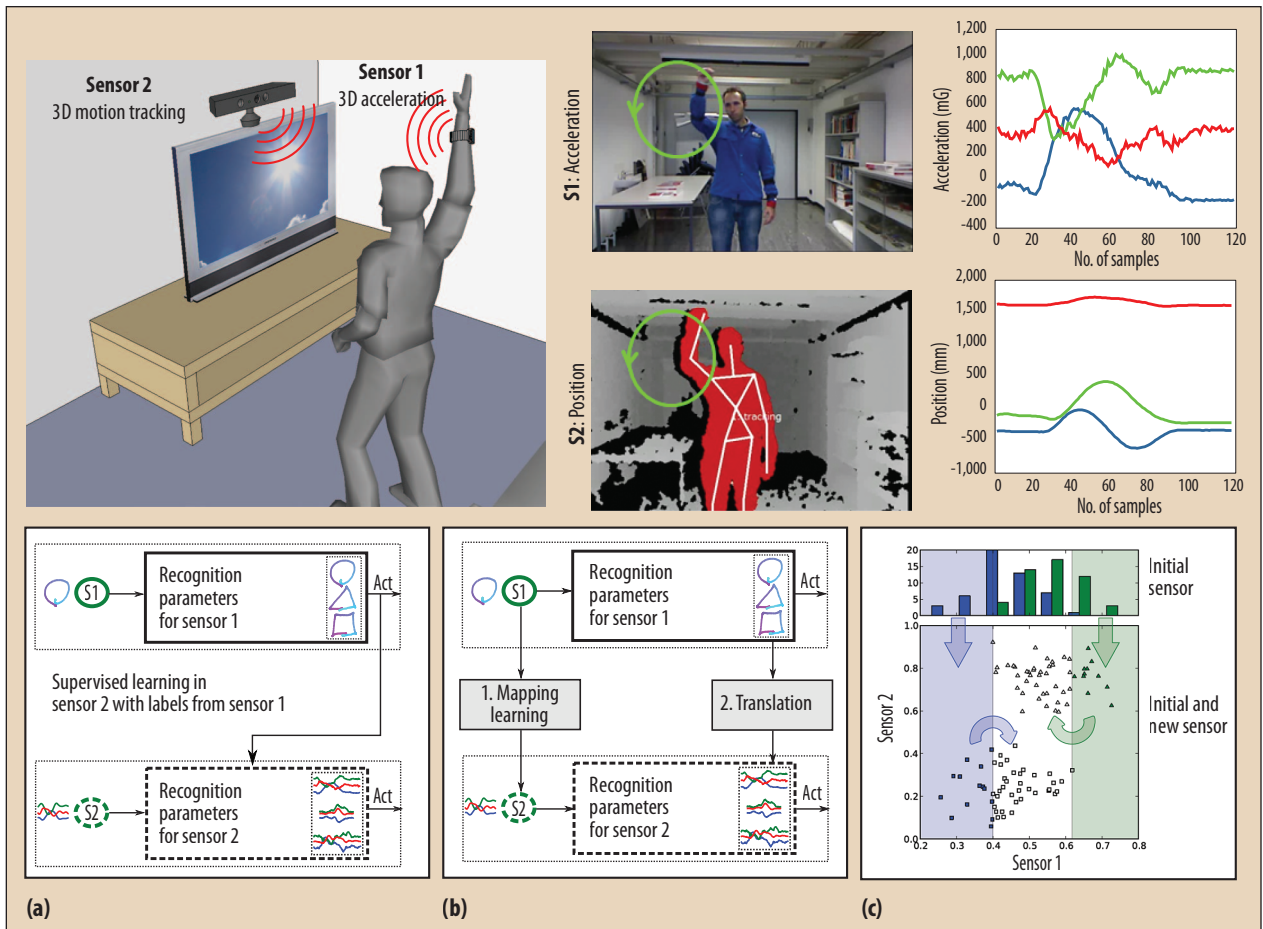
### Setting up the recognition goal

The scenario begins with the user returning home (1). To quantify his physical activity, we set the recognition goal to "locomotion." There are accelerometers in the user's smartphone, wristwatch, and smart clothing, but none of the nodes' ExperienceItems recognize this activity. After goal reasoning, the system determines that "walk" and "run" are aspects of locomotion.

The smartphone indicates that it can recognize these activities when placed in a hip holster or trousers pocket, with the node's ExperienceItems specifying different recognition-chain parameters for these locations. We found that a device can autonomously identify the body location at which it is being carried by analyzing its movement patterns.<sup>8</sup> The system thus can instantiate a recognition chain for locomotion according to the phone's position. If the phone changes location, the system updates the node's self-description and reconfigures itself.

### Substituting sensors

The user enters his home and heads to the living room to relax, leaving his smartphone on a table (2). Although the system can no longer use the phone to perform activity recognition, the user is within range of a passive infrared (PIR) motion sensor for lighting control. The PIR sensor rec-



**Figure 3.** Using an unknown new sensor node (wristwatch accelerometer) that coexists with a preexisting node (Microsoft Kinect for Xbox 360 motion-capture device) for opportunistic activity recognition. (a) The preexisting node supervises incremental learning of the new node when it detects an activity or context. (b) System identification makes it possible to map the two nodes' signals. The mapping function translates the first node's recognition chain for use by the second node. In both (a) and (b), the new node learns to recognize the same events as the preexisting node, which can disappear. (c) Activities that are confused in the feature space spanned by one sensor (top) can be disambiguated by expanding the feature space with the second sensor's dimensions (bottom). The two nodes must coexist.

ognizes “movement,” and the system uses goal reasoning to determine that this constitutes locomotion in the house. The recognition chain instantiated using the PIR sensor does not need complex features or classifiers: movement simply means locomotion.

### Exploiting unknown sensors with transfer learning

Unknown new sensors will appear in the IoT with infrastructure upgrades, or when the user buys new gadgets. These sensors might not self-describe if they are legacy devices. Nevertheless, being able to use an unknown sensor automatically would reduce system maintenance and programming costs. Our approach thus uses an autonomous process that “spreads” recognition capabilities to new nodes, much like AI agents teaching each other tasks.

*Transfer learning* consists of translating the knowledge available to solve a problem in one domain to a different

but related domain. We devised a lightweight implementation of transfer learning to translate the ability to recognize activities from one node to another.

As Figure 3a shows, when a preexisting node detects an activity, it triggers an incremental supervised learning process for task recognition on a new node. The system estimates the new node's performance by comparing its decisions to those of the preexisting node. If transfer learning is successful, a new ExperienceItem is created on the new node, which can now help fulfill the recognition goal. We have used this approach to transfer recognition tasks for locomotion and body postures (stand, sit, walk, lie) to eight different on-body accelerometers.<sup>9</sup>

Continuing with our scenario, the user enters the kitchen to prepare lunch (3). His wristwatch can sense arm movements with an acceleration sensor, but this node has no relevant ExperienceItems to identify these movements as locomotion. However, people also typically move their



arms when walking or running. In this case, a PIR sensor in the kitchen detects the user's movement and then teaches this recognition task to the user's wristwatch.

### Exploiting behavioral assumption

Some consumer appliances, such as clothes dryers and ovens, have sensors indicating that the door is open to stop operation. These sensors can also provide information about a user's activities when combined with assumptions about human behavior. For example, when the appliance detects the door opening or closing, it can indicate an "open" or "close" gesture by the user and provide a new ExperienceItem to the wristwatch.<sup>9</sup>

Behavioral assumptions can also be applied to locomotion: a user is likely to "stand" when opening a door and afterward will "walk" away. The Opportunity framework can apply numerous behavioral assumptions derived from the large annotated dataset of activities in the UCI Machine Learning Repository (<http://archive.ics.uci.edu/ml/datasets/opportunity+activity+recognition>).

### Exploiting unknown sensors through system identification

After lunch, the user returns to the living room to play a game using his Microsoft Kinect system (4). The Kinect is a depth camera with motion-tracking as well as facial- and voice-recognition capabilities: it delivers 3D coordinates of the user's joints to the Xbox 360 console, enabling him to simply move his body to control an on-screen avatar.

The opportunistic activity recognition system can relate this data to that of other unknown sensors, including the user's wristwatch accelerometer, using various system identification techniques. As Figure 3b shows, the system finds this mapping automatically at runtime, translating the Kinect's recognition chain for use by the wristwatch

---

## The node ensembles that the IoT makes available offer a promising way for an opportunistic activity recognition system to utilize multiple "critics."

---

sensor. Using system identification, we have successfully translated the capacity to recognize five hand gestures between the Kinect and body-worn accelerometers with performance degradation of less than 4 percent in as short a time as 3 seconds.<sup>10</sup> Many other sensor transformations could be achieved in this way.

### Exploiting unknown sensors for unsupervised disambiguation

We have developed a method, inspired by semisupervised learning techniques, to use a newly discovered sensor

without ExperienceItems to improve a known sensor's accuracy. The top part of Figure 3c shows two activities that overlap in a feature space spanned by one sensor and thus cannot be recognized without errors. As the bottom part of Figure 3c shows, expanding the feature space with the new sensor's dimensions helps disambiguate the two activities, which now form distinct clusters. A systematic analysis of three sets of data and thousands of pairs of sensors showed that it is possible to enable this method only when disambiguation is likely to be beneficial.<sup>11</sup> Once it discovers this information, a sensor can store it in its ExperienceItems for future use.

### Self-adaptation

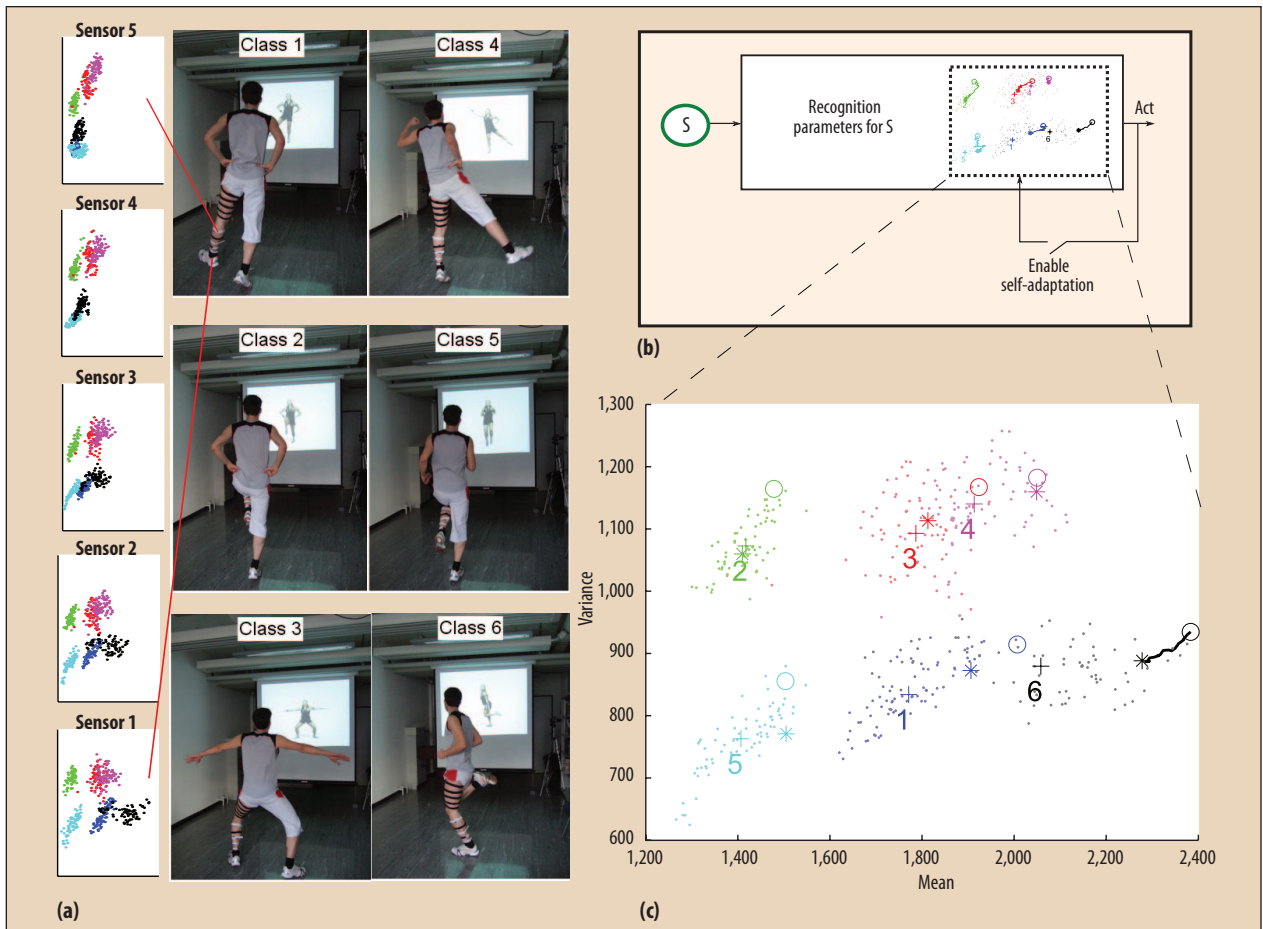
Our scenario concludes with the user, outfitted in smart clothing, going jogging (5). As Figure 4 shows, intense physical activity can lead to sensor displacement, which in turn causes the sensor signal patterns to change. To address this problem, we developed an expectation-maximization method using a nearest class-center classifier that enables the opportunistic activity recognition system to dynamically self-adapt to the variable statistics of sensor signals.<sup>9,12</sup> This technique improved performance when sensors were displaced along a limb. Related methods are applicable to other moderate changes, such as sensor rotation or sensitivity change.

### Exploiting ensembles to assess performance

An autonomous system must assess its own performance. It usually accomplishes this via a "critic"—a module that estimates performance from signals available within the system. The node ensembles that the IoT makes available offer a promising way for an opportunistic activity recognition system to utilize multiple critics. Because the nodes in an ensemble work toward the same goal, the ensemble can assess the performance of its individual nodes, or control their adaptation, as long as most of the nodes fulfill their mission. We have developed several methods to exploit this capability.

Harsh environments or strenuous activities can damage sensors—for example, fibers in smart clothing can rupture while the user is exercising. A critic can detect anomalous sensors in a node ensemble by looking at the statistical agreement between the nodes over time. Significant changes in the agreement patterns indicate an anomaly. Upon detecting such an anomaly, the system either adapts the dynamic ensemble to avoid the faulty node or it includes additional nodes.

There is, however, no way to estimate the performance of a fused set of nodes based solely on the accuracy parameters in their ExperienceItems. Indeed, this requires a pairwise measure of their complementarity, known as diversity,<sup>12</sup> and precomputing diversity is not appropriate in an opportunistic setting. For this reason, the system uses



**Figure 4. Self-adaptation.** (a) Sensor placement affects the mapping between activity classes and features derived from sensors, in this case for six fitness activities. (b) Self-adaptation makes it possible to incrementally adjust the activity models to follow the distributions of the classes in the feature space. (c) Adaptation dynamics of a nearest class-center classifier trained on one sensor position and deployed on another position.

another critic to estimate diversity at runtime by observing the decision agreement between individual classifiers and the fused decision. It then uses an information-theoretic approach to build classifier ensembles based on the diversity estimated by the critic and the accuracy reported in the nodes' ExperienceItems. Experimental results showed that this approach leads to high-performance ensembles.<sup>12</sup>

## LESSONS LEARNED

The opportunistic activity recognition paradigm is ideally matched to the big data that the IoT makes available and could help fulfill the promise of ambient intelligence anywhere, anytime, using a plethora of readily deployed sensor nodes.

Our work addresses the gap between low-level activity recognition methods, which require training data from predefined sensor configurations, and higher-level contextual abstractions, which assume the existence of services that can be semantically replaced or combined. Functionally, our approach is complementary to SOAs and can

be used to create reliable applications despite changing sensor availability. The opportunistic principles can be included as part of the firmware of smart things or within an overarching IoT framework. As IoT-enabled “things” become more common, an opportunistic activity recognition system becomes more capable as it can exploit larger node ensembles and discover more elaborate solutions to a recognition problem. Most methods are equally applicable to ambient, object, or wearable sensors.

Our approach tightly integrates signal processing and machine learning techniques to identify data patterns across multiple noisy sensors with a contextual framework to handle goals, sensor capabilities, and sensor transformations and to reason about complex activities. For instance, the system will instantiate an opportunistic method to carry out transfer learning between specific sensor modalities only in appropriate situations. In addition, the system includes runtime capabilities that are important to discovering new solutions to the recognition problem and to operating in open-ended environments.





This integration of low- and high-level approaches enables us to translate AI-based principles of autonomous operation into a practical ambient intelligence system.

If opportunistic activity recognition is to become a mainstream technology, commercial vendors and researchers must work together to create adequate sensor self-descriptions and activity ontologies for future smart things. In addition, there is a need for standardized methodological evaluations and performance benchmarks for recognition systems. We must also assess the stability-plasticity tradeoff to avoid catastrophic failures in long-term deployments. This poses new design and validation challenges that might be addressed using control theory.

In the near term, opportunistic activity recognition brings immediate benefits to the IoT:

- the reuse of readily deployed sensors will reduce system costs;
- the ability to roam around our surroundings and seamlessly substitute various sensors will improve user acceptance;
- reconfiguring systems to mitigate sensor losses will increase robustness; and
- systems will autonomously expand to include newly deployed resources, simplifying scale-up.

Decoupling pervasive application design from sensor deployment could create a market that is as hot as the current smartphone apps market. The decreased reliance on dedicated infrastructure will make the IoT accessible to enthusiasts as well as established players, leading to the creative exploration of many novel concepts in ambient intelligence.

Looking farther down the road, ambient intelligence systems could learn to discover autonomously when and how to use actuators accessed through the IoT to maximize user satisfaction. Achieving this next level of intelligence will require more human-like machine learning techniques<sup>13</sup> and implicit understanding of users' needs.<sup>14</sup> **C**

## Acknowledgments

The Opportunity project acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open grant number 225938. The authors acknowledge the contributions of their Opportunity colleagues: David Bannach, Alberto Calatroni, Kilian Förster, Gerold Hölzl, Marc Kurz, and Hesam Sagha.

## References

1. M. Weiser, "Ubiquitous Computing," *Computer*, Oct. 1993, pp. 71-72.
2. M. Presser et al., "The SENSEI Project: Integrating the

Physical World with the Digital World of the Network of the Future," *IEEE Comm. Magazine*, vol. 47, no. 4, 2009, pp. 1-4.

3. D.J. Cook, J.C. Augusto, and V.R. Jakkula, "Ambient Intelligence: Technologies, Applications, and Opportunities," *Pervasive and Mobile Computing*, Aug. 2009, pp. 277-298.
4. A. Rezgui and M. Eltoweissy, "Service-Oriented Sensor-Actuator Networks: Promises, Challenges, and the Road Ahead," *Computer Comm.*, Sept. 2007, pp. 2627-2648.
5. J.O. Kephart and D.M. Chess, "The Vision of Autonomic Computing," *Computer*, Jan. 2003, pp. 41-50.
6. M. Conti and M. Kumar, "Opportunities in Opportunistic Computing," *Computer*, Jan. 2010, pp. 42-50.
7. K. Fuji and T. Suda, "Semantics-Based Context-Aware Dynamic Service Composition," *ACM Trans. Autonomous and Adaptive Systems*, May 2009, article no. 12.
8. K. Kunze and P. Lukowicz, "Symbolic Object Localization through Active Sampling of Acceleration and Sound Signatures," *Proc. 9th Int'l Conf. Ubiquitous Computing (UbiComp 07)*, Springer, 2007, pp. 163-180.
9. D. Roggen et al., "Wearable Computing: Designing and Sharing Activity-Recognition Systems across Platforms," *IEEE Robotics & Automation Magazine*, vol. 18, no. 2, 2011, pp. 83-95.
10. O. Baños et al., "Kinect=IMU? Learning MIMO Signal Mappings to Automatically Translate Activity Recognition Systems across Sensor Modalities," *Proc. 16th Int'l Symp. Wearable Computers (ISWC 12)*, IEEE, 2012, pp. 92-99.
11. D. Bannach, B. Sick, and P. Lukowicz, "Automatic Adaptation of Mobile Activity Recognition Systems to New Sensors," *Proc. 2011 Workshop Mobile Sensing: Challenges, Opportunities, and Future Directions (UbiComp 11)*, ACM, 2011; [http://research.microsoft.com/en-us/um/beijing/events/ms\\_ubicomp11/papers/bannach.pdf](http://research.microsoft.com/en-us/um/beijing/events/ms_ubicomp11/papers/bannach.pdf).
12. R. Chavarriaga, H. Sagha, and J. del R. Millán, "Ensemble Creation and Reconfiguration for Activity Recognition: An Information Theoretic Approach," *Proc. IEEE Int'l Conf. Systems, Man, and Cybernetics (IEEE SMC 11)*, IEEE, 2011, pp. 2761-2766.
13. J.B. Tenenbaum et al., "How to Grow a Mind: Statistics, Structure, and Abstraction," *Science*, Mar. 2011, pp. 1279-1285.
14. M. Pantic and A. Vinciarelli, "Implicit Human-Centered Tagging," *IEEE Signal Processing Magazine*, vol. 26, no. 6, 2009, pp. 173-180.

*Daniel Roggen is a senior researcher at the Wearable Computing Laboratory, ETH Zurich, Switzerland. His research interests include ubiquitous computing, activity recognition, and robotics. He is coordinator of the EU FP7 FET-Open Opportunity project. Roggen received a PhD in microengineering from École Polytechnique Fédérale de Lausanne, Switzerland. He is an IEEE associate member. Contact him at droggen@gmail.com.*

*Gerhard Tröster is head of the Wearable Computing Laboratory at ETH Zurich, Switzerland. His research interests include wearable computing for healthcare and production,*

smart textiles, and flexible electronics. Tröster received a PhD in electrical engineering from Darmstadt University of Technology, Germany. Contact him at [troester@ife.ee.ethz.ch](mailto:troester@ife.ee.ethz.ch).

**Paul Lukowicz** is a professor at DFKI and Kaiserslautern University of Technology, Germany. His research interests include wearable and mobile computer architecture, and context and activity recognition. Lukowicz received a PhD in computer science from the University of Karlsruhe, Germany. Contact him at [paul.lukowicz@dfki.de](mailto:paul.lukowicz@dfki.de).

**Alois Ferscha** is head of the Institute for Pervasive Computing at Johannes Kepler University of Linz, Austria. His research interests include pervasive and ubiquitous computing, networked embedded systems, embedded software systems, wireless communication, and distributed interactive simulation. Ferscha received a PhD in business informatics from the University of Vienna, Austria. Contact him at [ferscha@pervasive.jku.at](mailto:ferscha@pervasive.jku.at).

**José del R. Millán** holds the Defitech Foundation Chair in Non-Invasive Brain-Machine Interface and is an associate professor at École Polytechnique Fédérale de Lausanne, Switzerland. His research interests include the use of brain

signals for multimodal interaction, noninvasive brain-controlled robots, and neuroprostheses. Millán received a PhD in computer science from the Technical University of Catalonia, Spain. Contact him at [jose.millan@epfl.ch](mailto:jose.millan@epfl.ch).

**Ricardo Chavarriaga** is a senior researcher at École Polytechnique Fédérale de Lausanne, Switzerland. His research focuses on robust brain-machine interfaces and multimodal human-machine interaction. Chavarriaga received a PhD in computational neuroscience from École Polytechnique Fédérale de Lausanne. Contact him at [ricardo.chavarriaga@epfl.ch](mailto:ricardo.chavarriaga@epfl.ch).



Selected CS articles and columns are available for free at <http://ComputingNow.computer.org>.

computing **now**

**NEW+EXPANDED**

IEEE COMPUTER SOCIETY'S **COMPUTING NOW** WEBSITE

Learn industry solutions you can use from practical articles, case studies, blogs, and interviews that address high-interest, focused areas of technology.

Mobile Computing • Cloud Computing • Security • Software • High-Performance Computing • Networking



IEEE  computer society

Visit <http://computingnow.computer.org>