Multimodal Wearable Sensing for Fine-Grained Activity Recognition in Healthcare

State-of-the-art in-home activity recognition schemes with wearable devices are mostly capable of detecting coarse-grained activities (sitting, standing, walking, or lying down), but can't distinguish complex activities (sitting on the floor versus on the sofa or bed). Such schemes aren't effective for emerging critical healthcare applications — for example, in remote monitoring of patients with Alzheimer's disease, bulimia, or anorexia — because they require a more comprehensive, contextual, and fine-grained recognition of complex daily user activities. Here, a novel approach for in-home, fine-grained activity recognition uses multimodal wearable sensors on multiple body positions, along with lightly deployed Bluetooth beacons in the environment.

Debraj De, Pratool Bharti, Sajal K. Das, and Sriram Chellappan
Missouri University of Science and Technology

Considerable literature exists on recognizing users' activities of daily living (ADLs). But the broad subject of human activity recognition at home has a large variation, based on the complexity of detected activity, and privacy or deployment overhead of sensing resources. On the one hand, most existing works are able to recognize only basic coarse-grained ADLs with wearable or infrastructural sensing, without direct privacy concerns. On the other hand, camera or video imaging can be used for fine-grained ADLs or instrumental activities of daily living (IADLs) recognition.[1] However, this has significant privacy concerns for major applications such as remote, assessment-based, in-home healthcare applications. The ADLs are typically basic self-care skills that people learn during early childhood, such as sitting, standing, walking, and watching TV. But IADLs are more complex tasks needed for independent living, such as cooking, housekeeping, and doing laundry. In essence, the ADLs often include physical or postural activities, while IADLs include activities that require a combination of physical and cognitive capabilities.

In the literature, only limited work exists on recognition of fine-grained, complex ADLs and IADLs with minimal infrastructure overhead and without direct privacy concerns. Complex activity recognition using resource-constrained sensors is a challenging

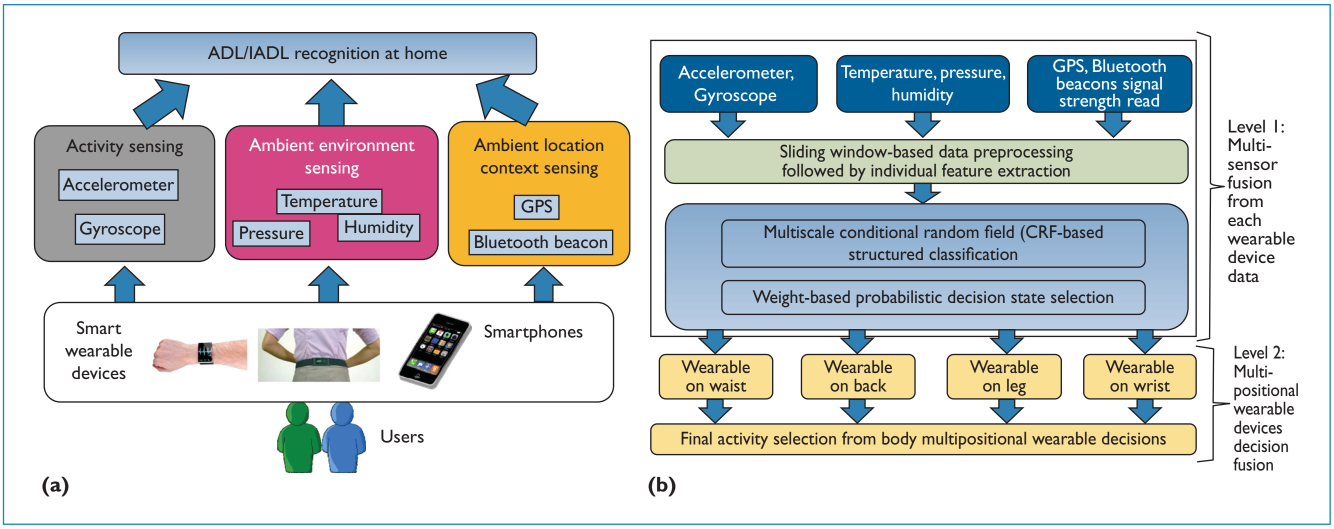

Figure 1. Our proposed system. (a) Multimodal sensing used in the proposed activities of daily living (ADL)/instrumental ADL (IADL)recognition system. (b) Our proposed activity classification algorithm from multimodal and multipositional sensing on a user's body.

task and requires innovative solutions. To this end, we propose an algorithm for complex ADL/IADL recognition that has potential application in healthcare and wellness management, such as remotely monitoring the progression of Alzheimer's disease in the elderly.

Elderly people affected by Alzheimer's disease (the most common form of dementia) often exhibit forgetfulness, memory loss, and repetitive behavior such as wandering outdoors at unusual hours and frequencies, wandering indoors in confusion, oversleeping at daytime, repetitive opening/closing of cabinet doors, and increased chances of falling.[2] Healthcare specialists (such as doctors, nurses, or caregivers/relatives) can use our proposed ADL/IADL activity-recognition scheme to remotely monitor and assess cognitive health progression (improvement, stability at the current stage, or degradation) in elderly people, while they're at home. Such cognitive health assessment could be critical for doctors in deciding whether it's time for a patient to be moved to assisted living or other formal care facilities. In health conditions such as bulimia (recurrent and frequent episodes of eating) and anorexia (severe restrictions on food intake), the frequencies and durations of eating are anomalous,[3] implying the need for remote monitoring and assessment of the degree of the eating disorder. Professionals can use our ADL/IADL activity recognition scheme not only to detect eating habits, but also to understand which other activities lead to or follow those eating phases.

Figure 1a illustrates our proposed ADL/IADL recognition system that exploits multimodal, wearable sensing of key contexts: *body locomotion* (via accelerometer and gyroscope), *ambient environment* (via temperature and humidity sensors, and fine-grained change in altitude of user body position via barometric pressure sensor), and *location* (via communication with Bluetooth beacon location tags). As for the infrastructure component, one simple, small, inexpensive Bluetooth beacon device is placed in each room. These beacons just broadcast, while user's wearable devices listen to them for assessing the location/proximity context. The activity classification algorithm (see Figure 1b) consists of three components:

- extraction of practical features from multimodal sensor suites on the wearable;
- a novel structured classification algorithm, based on the conditional random field (CRF) classifier,[4] performed on each of the wearable devices (on multiple positions of the body) separately; and
- improvement in activity classification by leveraging and fusing decisions from the body's multipositional wearable devices.

Through real experimental study we validated that the proposed scheme can classify a set of 19 complex, in-home activities with a high degree of accuracy overall. These activities are

Related Works in Activity Recognition

Over the years, many have researched human activity recognition, generating work that provides some background on the state of the field. Here, we discuss the most relevant of such works.

Oscar Laras and Miguel Labrador provided a detailed survey on human activity recognition schemes based on wearable sensors.[1] In another work, Saguna Saguna and her colleagues discuss complex activity recognition theory.[2] Based on sensing resources, there are three broad categories of approaches proposed in the literature for in-home activity recognition:

1. only with wearable devices,[3-5]
2. only with external static infrastructure-based systems deployed in surrounding physical environments,[6] and
3. combining wearable devices and external static infrastructure-based systems.[7]

With wearable devices only, activity recognition can be accomplished with the help of learning from data sensed by smartphones, wearable health tracker devices, smartwatches, augmented reality devices (such as Google Glass), and so on. For example, Kai Zhan and his colleagues used a smartphone-based accelerometer and first-person-view video camera for recognizing both locomotive and stationary activities.[3]

For static infrastructure-based systems, we can achieve activity recognition by learning from data sensed by sensor networks, telemetry systems, motion detectors, RFID tags and readers, video and image cameras, or smart appliances deployed in physical environments. For example, in Daniel Wilson and Chris Atkeson's work, they used static infrastructure-based sensors for room-level tracking and basic activity recognition such as sleeping in a bed, user movement status, and so on.[6]

Finally, there also exists a line of work that achieves activities of daily living (ADL) recognition by combining data sensed from wearable devices and additional static infrastructure. As an example, Nirmalya Roy and his colleagues' work integrates data sensed from networked motion sensors (that also provides location contexts) with those generated from smartphone sensors for classifying postural/locomotive activity states of multiple inhabitants.[7]

Table A provides a comparison of relevant existing works with our proposed ADL/instrumental ADL (IADL) recognition scheme. Note that although there's a large body of literature concerning a variety of activity recognition, here we focus on relevant works that detect only in-home activities and don't use sensing modes with direct privacy concerns. In particular, our approach uses very light additional infrastructure (Bluetooth beacon tags in each room) and has no direct privacy concerns (regarding infrastructure or a wearable camera), yet recognizing a larger number of fine-grained complex activities of individuals at home.

References

1. O. Lara and M. Labrador, "A Survey on Human Activity Recognition Using Wearable Sensors," *IEEE Comm. Surveys Tutorials*, vol. 15, no. 3, 2013, pp. 1192–1209.
2. S. Saguna, A. Zaslavsky, and D. Chakraborty, "Complex Activity Recognition Using Context-Driven Activity Theory and Activity Signatures," *ACM Trans. Computer-Human Interaction*, vol. 20, no. 6, 2013, pp. 32:1–32:34.
3. K. Zhan, S. Faux, and F. Ramos, "Multi-Scale Conditional Random Fields for First-Person Activity Recognition," *Proc. IEEE Int'l Conf. Pervasive Computing and Comm.*, 2014, pp. 51–59.
4. L. Wang et al., "Recognizing Multi-User Activities Using Wearable Sensors in a Smart Home," *Pervasive and Mobile Computing*, vol. 7, no. 3, 2011, pp. 287–298.
5. Z. Zhao et al., "Cross-People Mobile Phone-Based Activity Recognition," *Proc. 22nd Int'l Joint Conf. Artificial Intelligence*, vol. 3, 2011, pp. 2545–2550.
6. D.H. Wilson and C. Atkeson, "Simultaneous Tracking and Activity Recognition (Star) Using Many Anonymous, Binary Sensors," *Proc. 3rd Int'l Conf. Pervasive Computing*, 2005, pp. 62–79.

- walk and run indoors (categorized as locomotive activities);
- use refrigerator, clean utensil, cook, sit and eat, and use bathroom sink (categorized as semantic activities, because they are related to certain routines or tasks belonging mostly to IADLs);
- move from indoor to outdoor, move from outdoor to indoor, walk upstairs, and walk downstairs (categorized as transitional activities); and
- stand, lie on the bed, sit on the bed, lie on the floor, sit on the floor, lie on the sofa, sit on the sofa, and sit on the toilet (categorized as postural/relatively stationary activities).

To the best of our knowledge, this is the first effort to recognize 19 in-home activities mainly with wearable devices. This is in contrast to typically 6 to 12 in-home activities recognized in existing literature,[5-7] as Table A shows (see the "Related Works in Activity Recognition" sidebar).

Complex ADL/IADL Recognition Scheme

We now describe the underlying methodology of our proposed in-home ADL/IADL recognition system (see Figure 1b). It consists of the following pipeline of three phases:

Table A. Comparison of different categories of works related to in-home activities of daily living (ADL)/ instrumental ADL (IADL) recognition.

Referenced work	Wearable sensing component	Infrastructural sensing component	Activities recognized
Daniel Wilson and Chris Atkeson[6]	None	Motion detectors, break-beam sensors, pressure mats, and contact switches	Room-level tracking and basic activities such as sleeping in bed, and user-movement status.
Piyush Gupta and colleagues[8]	Belt-clip accelerometer	None	Six activities: walking, jumping, running, sit-to-stand/stand-to-sit, stand-to-kneel-to-stand, and staying stationary.
Kai Zhan and colleagues[3]	Accelerometer Wearable video camera	None	Twelve activities: walking, going upstairs, going downstairs, drinking, stand up, sit down, sitting, reading, watching TV/monitor, writing, turning water faucet on/off, and hand washing.
Nirmalya Roy and colleagues[7]	Smartphone accelerometer and gyroscope	Ceiling-mounted infrared motion sensors	Six low-level postural or motion activities (sitting, standing, walking, running, lying and climbing stairs), and six high-level semantic activities (cleaning, cooking, medication, sweeping, washing hands, and watering plants).
Complex ADL/ IADL recognition (this article)	Wearable multimodal sensing with • accelerometer, gyroscope (for body locomotion); • temperature, atmospheric pressure, and humidity (for ambient environment); and • Bluetooth message reception (for location context).	Bluetooth beacons in the physical environment (one beacon in each room)	Nineteen fine-grained activities: • walking and running indoors (locomotive); • using refrigerator, cleaning utensils, cooking, sitting and eating, and using bathroom sink (semantic); • indoor to outdoor, outdoor to indoor, and walk upstairs or downstairs (transitional); and • just standing, lying on bed, sitting on bed, lying on floor, sitting on floor, lying on sofa, sitting on sofa, and sitting on toilet (postural/ stationary).

7. N. Roy, A. Misra, and D. Cook, "Infrastructure-Assisted Smartphone Based ADL Recognition in Multi-Inhabitant Smart Environments, *Proc. IEEE Int'l Conf. Pervasive Computing and Comm.*, 2013, pp. 38–46.
8. P. Gupta and T. Dallas, "Feature Selection and Activity Recognition System Using a Single Triaxial Accelerometer," *IEEE Trans. Biomedical Eng.*, vol. 61, no. 6, 2014, pp. 1780–1786.

1. data preprocessing followed by feature extraction on each of the sensor data stream (which runs separately on each wearable device's data);
2. multiscale CRF-based classification,[4] followed by the proposed weight-based probabilistic decision state selection (which runs separately on each wearable device data); and
3. final user activity state classification, given an individual decision from each multipositional wearable device on the body.

The sampled multimodal sensor data are individually preprocessed and fed to the feature-extraction process. The extracted features are then used in the multiscale CRF classifier for supervised learning. But instead of using the deterministically best decision about the activity state, our modified classifier uses a weight-based, probabilistic activity state selection. This selection is done from the set of top K classified activities and their emission/output probabilities. Finally, the classifier decisions from individual wearable devices are converted into a final activity state using a body-position-based selection. This last phase of decision selection is flexible, based on the number of simultaneous wearable devices worn by the user.

With growing commercial use and development of smart wearable devices, it won't be uncommon to see users with multiple wearable devices on different body locations. In fact, there already are popular wearable devices on the market, such as Lumo Back (on the waist or lower back; see www.lumobodytech.com/lumoback), Lumo Lift (on the back; see www.lumobodytech.com/lumolift), Nike+ (on legs or shoes; see https://secure-nikeplus.nike.com/plus), and Fitbit (on the wrist; see https://www.fitbit.com). Upcoming wearable health trackers are going to host not only the movement- and posture-based sensors, but also ambient environment sensors, and also will record communication signatures with Bluetooth beacons in the surrounding infrastructure. This shows potentially broad practical applicability of our designed system and algorithm.

Next, we describe each of the three pipeline phases in detail.

Phase I: Feature Set Extraction

Computation of meaningful features from the raw multimodal sensor data is important. To extract the feature dataset, we employ a selection of features that are lightweight to compute, moderate in number, and robust to practical situations that might affect their effectiveness (for example, changing the tile or rotation of the wearable).

Accelerometer

Features from *accelerometer* data are collected as follows. Several works in the literature propose to compute too many features from the accelerometer data streams, which can often be time- and resource-consuming due to the computational load and fast sampling required for an accelerometer. Realizing this, and after analyzing the literature on accelerometer data processing, our proposed system uses six particular features from accelerometer data for each sampled sliding window (with the sliding window's size and duration generally set at two seconds): mean and variance of resultant acceleration $\left(\sqrt{a_x^2 + a_y^2 + a_z^2}\right)$, where a_x, a_y, and a_z are accelerated along the x-, y-, and z-axes), with mean and variance of the first derivative of resultant acceleration, and mean and variance of the second derivative of resultant acceleration.

The resultant acceleration is a good measure of the degree of body movement due to activity. These features are commonly used in the related literature. Our system doesn't use as many features so that it can support online activity classification (computation of too many features will bring considerable delay).

More importantly, note that axis-specific parameters of the accelerometer aren't used for practical applicability in real-world scenarios. All six features are the combined property of all three axes, and thus not affected by continuous change in rotation or tilt of the wearable devices. The three-axis accelerometer is sampled at 100 Hz. This sampling frequency is satisfactory to capture acceleration-based human users' regular body movements.[5]

Gyroscope

Features from the *gyroscope* data are collected as follows. The proposed system uses the following six features from gyroscope data for each sampled sliding window (again with a size of two seconds): mean and variance of resultant (of the three-axis) angular speed, mean and variance of first derivatives of resultant angular speed, and mean and variance of the second derivative of resultant angular speed. These features from a gyroscope are also commonly used in the related literature. However, as previously mentioned, our system doesn't use extra features that would delay computation. Also, the axis-specific parameters of a gyroscope aren't used here for practical applicability in real-world scenarios. All six features are the combined property of all three axes. The three-axis gyroscope is sampled at 100 Hz.

Temperature, Humidity, and Atmospheric Pressure

Features from these sensors are collected as follows. They're sampled at 1 Hz, 1 Hz, and 5 Hz, respectively. The system used the mean and variance of windowed data for each of these sensors (again with the sliding windows' size typically chosen as two seconds). The pressure sensor sampling rate is kept a little higher to capture fine changes in atmospheric pressure in different user positions at slightly different heights (elevations), such as on the floor, sofa, bed, or chair. We observed that the barometric pressure sensor indicates subtle differences in pressure data at those elevations. Although the differences are subtle, it can be exploited through machine-learning-based supervised

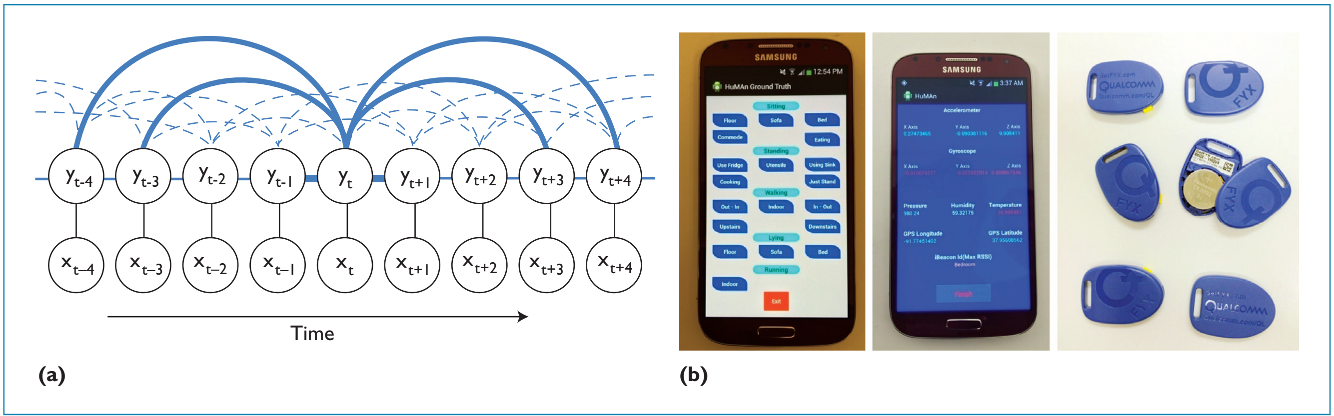

Figure 2. The conditional random field (CRF) model. (a) Multiscale CRF graph structure. The thick edges represent the pairwise edges for the template setting of (010304) for hidden state node y_t. (b) From left to right: application screenshot for ground-truth data collector, smartphone application screenshot for sensor sampling, and Bluetooth beacons.

classification for successfully distinguishing complex activities. While moving upstairs or downstairs, the atmospheric air pressure change is distinct. Also, subtle changes of temperature and humidity around wearables on certain body positions help detect and distinguish complex and fine-grained activities. For example, the temperature and humidity around the wrist-worn wearable changes during activities such as opening the refrigerator, using the sink, cooking, transitioning from indoors to outdoors, and so on. This motivates us to monitor the user's sensing of physical ambience surrounding temperature, humidity, and atmospheric pressure.

Bluetooth Beacons

Another novel feature of the proposed system is determining and using the location context through message reception on a wearable device from Bluetooth beacon broadcasts. These Bluetooth beacon devices (see Figure 2b) are becoming popular due to their low overhead in deployment and management. These are small transmitters that can notify nearby devices of their presence, representing proximity of those devices to the beacons. We exploited this practical feasibility to provide our system with location context-based features. The system uses a received signal strength indicator (RSSI) of the received message on the wearable device to estimate location context in its closest proximity. Because we equipped each room with one Bluetooth beacon, the features were basically an indicator of user presence in rooms, such as the bedroom, kitchen, or bathroom. However, in future work we aim to add more location features, and also refine location context detection with more analysis on RSSI data from beacon broadcasts.

It's worth mentioning that instead of dedicated Bluetooth beacons, someone also could use existing in-home Bluetooth devices (for example, from smart appliances or connected utility devices) as alternative sources of location contexts. Our proposed scheme can work with those already in-use Bluetooth devices by simply changing the source of location tags listed by the wearable Bluetooth receivers. The Bluetooth hardware ID to location tag mapping is flexible and used in the software that runs in the server, thus requiring no reconfiguration or reprogramming of devices.

Phase 2: Structured Classification of Activities

We based our proposed activity classifier on the CRF multiscale graphical model-based structured classifier.[4] It's a class of statistical modeling method used for structured learning and prediction. CRF can support more complex and useful feature sets by modeling the posterior probabilities instead of joint probabilities. Unlike ordinary classifiers, which predict activities for a single sample (from a combination of chosen features) without using neighboring samples, the CRF classifier can take into account the multiscale context. It's used in this work for modeling the multiscale context in activity sequences.

The CRF model captures the temporal relationships in sequential activity data. Figure 2a shows the graph structure used in the CRF model, with observation sequence **x** (obtained from the feature extraction of the multimodal sensor data), hidden states **y** (the user activity states to be classified) of class probability assignments, and the edges E between hidden states that represent pair-wise relationships. As in Figure 2a, the different scale or length of edges (for example, that run from y_t to y_{t-1}, y_{t+1}, y_{t-3}, y_{t+3}, y_{t-4}, y_{t+4}) enables flow of contextual information in the whole network. To implement the CRF model, we used a modified version of the standard CRFSharp toolkit (see http://crfsharp.codeplex.com).

Phase 3: Body Multipositional Decision Selection

The classification output from the modified CRF-based classifier consists of tuples: *<detected activity, body position of wearable>*. The number of such tuples is decided by the number of wearable devices needed — for example, four if waist, back, leg, and wrist positions are used. The final task is to select one overall user activity state from these pairs. We handle this selection as follows: we select each possible activity state to be most relevant to one of the body positions (for example, walking/running for leg position, lying down for back position, and so on). We base this relevance relationship on basic knowledge of human physiology and also on our experimental observations. Then the system makes the final decision about the user's activity state from the most relevant pair available. If there's a tie, the system chooses the final activity state randomly from the equally probable choice of pairs.

Note that this multipositional decision-selection methodology might not be the best solution. This is kept as future work, as the main focus of this work has been to first solve multimodal sensing-based classification. Also for future work, we aim to improve the multipositional device decision selection with another layer of machine-learning classifiers, which don't need a temporal order for learning.

Experimental Evaluation

Now we present preliminary experimental results to evaluate our ADL/IADL classification system's performance. We tested the classification performance with a user in a real-home setting. Our experimental setup includes Samsung Galaxy S4 smartphones' onboard sensors, along with Gimbal Bluetooth beacons (see www.gimbal.com). We used collected data from four smartphones worn on different body positions: waist, lower back, thigh, and wrist.

It's important to note that smartphones are used only as a multisensor data-collection platform. The wearable devices on different body positions will ideally be designed with different form factors in practical scenarios. Our next stage of work will include the design of wearable prototypes with those multimodal sensing capabilities. The available wearable devices on the market mostly don't have such specific sensors needed by our activity classifier and detector. For future work, we plan to use the RFduino-based small form-factor sensor platform and off-the-shelf sensing modules, to design the wearable prototype for further system integration and application.

We developed two Android applications to serve two purposes. First, to collect data from on-board selected sensors and receive signals from Bluetooth beacons installed in different rooms of the home environment (collected data was locally stored with proper timestamps; a clock of each device was Network Time Protocol [NTP]-time synchronized with an available Android app called NTPSync). Second, the Android applications collected user activity ground truth with proper timestamps (we also used a separate person as a dedicated observer for using the ground-truth application).

The user performed all 19 activities over the course of 45 minutes in the user's own selected order of choice and repetition (to keep the natural activity behavior, movements, and postures). The user did this twice to obtain separate training and testing datasets. Figure 3 shows the order of performed activities (the ground truth through time) in the testing dataset collection. We evaluated the proposed ADL/IADL classifier performance on the same user's activity data, with separate datasets for training and testing. Figures 4a and 4b show high classification accuracy performance. In our next stage of work, we aim to evaluate cross-user activity classification performance. In this work, however, our proposed scheme has shown to be highly effective for *personalized application* that requires fine-grained and complex activity recognition.

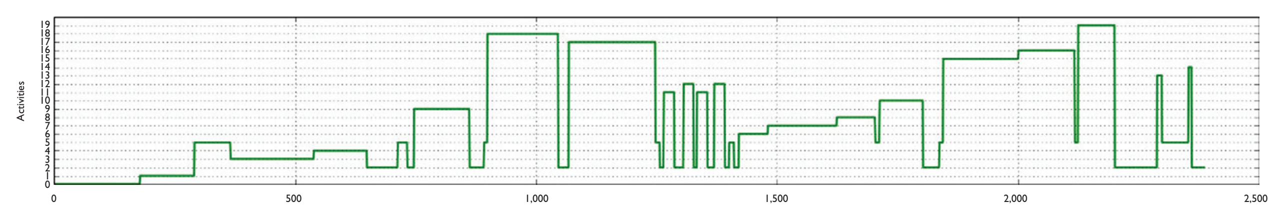

Figure 3. Ground truth. The order of 19 activities performed by the user (this shows the activity index versus sample number or time).

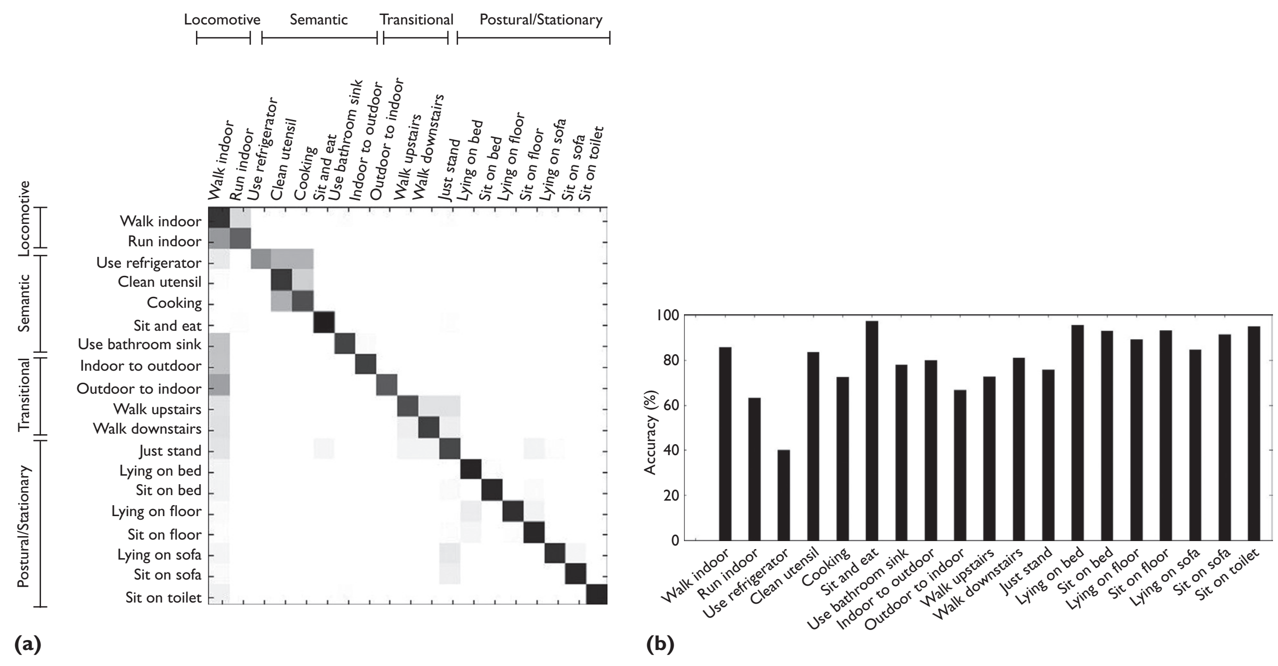

Figure 4. Same-user evaluation (with different datasets for training and testing). (a) A confusion matrix of accuracy performance across 19 activities. (b) Individual accuracy performance of 19 activities. From Figure 4b, we observe that the mean accuracy over all activities is 80.48 percent, and the median of accuracy over activities is 82.14 percent. The accuracy of individual activities was a minimum of 40 percent (using the refrigerator), and maximum of 97.3 percent (sitting and eating). Note that with the 19 complex activities classifier, a random guess will lead to an average of 100/19 percent or 5.26 percent accuracy. Thus, the accuracy percentage of 40 percent is still an improvement compared to the 5.26 percent accuracy of a random guess. Statistically, the requirement of 50 percent or higher accuracy isn't representative in this particular case, because this isn't a binary classifier or one with a much lower number of activities. Yet our proposed scheme achieved considerably higher accuracy for most of the 19 activities.

Furthermore, the accuracy errors for each ground truth of activity aren't randomly or equally distributed across all of the other remaining activities. The activity classification errors are distributed in only a small number of other similar activities. This is clearly observed in the confusion matrix shown in Figure 4a. Some complex/hybrid activities, such as opening the refrigerator, are more difficult to detect. One possible reason could be that detecting such activity partly relies on the temperature sensor on the wrist (which gives a unique temperature signature when a wearable device on the wrist is in close proximity to an opened refrigerator). Because the temperature sensor had some delay in sensing the

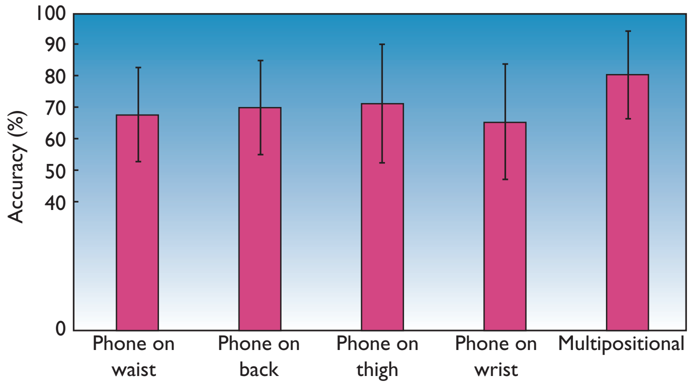

Figure 5. Average accuracy of multipositional decision making for same-user evaluation, compared to individual activity decisions from each of the wearable devices placed on the user's different body positions. The proposed activity classifier is run separately on each wearable device's data.

environment, it sometimes confused activities with a similar context in the kitchen, such as cleaning utensils and cooking. From the confusion matrix in Figure 4a, we observe that our classification error for using the refrigerator didn't spread randomly among other activities, but was distributed in just two other kitchen-related activities.

Figure 5 presents the performance evaluation of multipositional decision making in the final phase of our proposed scheme. For same-user evaluation (same-user training and testing on different datasets), the individual wearable devices (worn on the user's different body positions) provide accuracy ranging from 65.83 to 71.48 percent, while our multipositional decision-selection strategy provides an overall accuracy of 80.48 percent. It's important to note that the main thesis of this research is to design an improved classifier for a large number of complex activities by using multimodal sensing. The next step is to design a good final decision selection methodology. We've used the basic notion of human physiology-based activity decision selection from devices worn on the waist, back, thigh, and wrist. But we can further improve system performance by using a more intelligent design for final activity selection, which is the next goal of our research. We also aim to use another classifier for final decision selection, which will use temporal activity information from previous time instances.

Here, we presented a scheme for recognizing humans' fine-grained, complex, in-home activities using wearable multimodal sensing and minimal infrastructure with only Bluetooth beacons. In particular, we designed a context-based activity classifier based on hybrid multimodal sensors, and body multipositional wearables. Our experimental results demonstrate that, on average, the proposed scheme provides more than 80 percent accuracy of classifying 19 in-home activities.

For future work, we plan to improve the body multipositional device decision-selection algorithm. We also aim to design a custom wearable prototype and apply the complete system for specific applications, such as remote assessment of elderly people with dementia while they live independently at home.

Acknowledgments

We thank the anonymous reviewers and guest editors for valuable suggestions that helped us improve this article's quality. The US National Science Foundation supported this work in part through grants 1404673, 1404677, 1254117, and 1205695. The authors thank Rakesh Kumar and Satyaki Roy for their help in the experimental study.

References

1. W.D. Spector et al., "The Hierarchical Relationship between Activities of Daily Living and Instrumental Activities of Daily Living," *J. Chronic Diseases*, vol. 40, no. 6, 1987, pp. 481–489.
2. M. Warner, *Complete Guide to Alzheimer's-Proofing Your Home*, Purdue Univ. Press, 2000.
3. H.W. Hoek and D. van Hoeken, "Review of the Prevalence and Incidence of Eating Disorders," *Int'l J. Eating Disorders*, vol. 34, no. 4, 2003, pp. 383–396.
4. J.D. Lafferty, A. McCallum, and F.C.N. Pereira, "Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data," *Proc. 18th Int'l Conf. Machine Learning*, 2001, pp. 282–289.
5. P. Gupta and T. Dallas, "Feature Selection and Activity Recognition System Using a Single Triaxial Accelerometer," *IEEE Trans. Biomedical Eng.*, vol. 61, no. 6, 2014, pp. 1780–1786.
6. K. Zhan, S. Faux, and F. Ramos, "Multi-Scale Conditional Random Fields for First-Person Activity Recognition," *Proc. IEEE Int'l Conf. Pervasive Computing and Comm.*, 2014, pp. 51–59.
7. N. Roy, A. Misra, and D. Cook, "Infrastructure-Assisted Smartphone Based ADL Recognition in Multi-Inhabitant Smart Environments," *Proc. IEEE Int'l Conf. Pervasive Computing and Comm.*, 2013, pp. 38–46.

Debraj De is a postdoctoral fellow in the Department of Computer Science, Missouri University of Science and Technology. His current research interests are in the areas of smart healthcare, smart environments, machine learning, and wireless sensor networks. De has a PhD in computer science from Georgia State University. Contact him at ded@mst.edu.

Pratool Bharti is a PhD student at the Department of Computer Science, Missouri University of Science and Technology. His current research work is in smart healthcare. Contact him at pbtn4@mst.edu.

Sajal K. Das is the Chair of Computer Science and Daniel St. Clair Endowed Chair at the Missouri University of Science and Technology. His current research interests include wireless sensor networks, smart healthcare, cyber-physical systems, mobile and pervasive computing, security and privacy, and social networks. Das has a PhD in computer science from the University of Central Florida. He's an IEEE Fellow. Contact him at sdas@mst.edu.

Sriram Chellappan is an associate professor in the Department of Computer Science at the Missouri University of Science and Technology. His current research interests include social and technological aspects in smart healthcare and cybersecurity. Chellappan has a PhD in computer science and engineering from the Ohio State University. Contact him at chellaps@mst.edu.

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