Sensor Placement Variations in Wearable Activity Recognition

How do placement variations in user-carried electronic appliances influence human action recognition? The authors categorize possible variations, present a systematic evaluation about their impact on human action recognition, and discuss ways to compensate for such variations.

Kai Kunze
Osaka Prefecture University

Paul Lukowicz
Technical University of Kaiserslautern Germany

We're currently witnessing an explosion in sensor-equipped, on-body consumer devices, from smartphones to Google Glass to fitness trackers. In general, the quality of the sensors embedded in these devices is similar to the ones typically found in the dedicated wearable sensing systems used in activity recognition research, but there's one major difference. With a few exceptions,[1,2] the bulk of research in human activity recognition to date assumes known fixed sensor locations, which are often carefully chosen to suit a particular application. In contrast, mobile consumer devices typically aren't firmly fixed to the body—they're placed in a pocket or bag, where they can shift around and rotate into different orientations. Even devices specifically designed for concrete placement such as glasses or bracelets might at times be carried in a pocket or bag.

Motivated by these considerations, recent research interest has focused on understanding the influence of device placement variations on activity recognition systems. However, so far, that research is fragmented, with most work focusing on individual, narrow problems and specialized methods. An activity recognition researcher trying to understand how to improve his or her system with respect to placement variations would need to work through a lot of papers with no overall conceptual framework. This article aims to change that by providing a systematic understanding of the type and impact of variations, describing a comprehensive set of methods to deal with the different variation types, and illustrating how such methods work and are designed and evaluated.

Understanding Placement Variations

Generally, we can distinguish between three types of device placement variations: on-body placement, within-body displacement, and orientation.

On-Body Placement

Depending on preference, people carry devices on different parts of their body. Variations based on these differences are usually coarse—for example, the device might reside in a front or back trouser pocket, a jacket pocket, an arm or a hip holster, or a bag.[3] Changes in on-body placement aren't as frequent as within-body displacements or changes in orientation—but even though several devices are associated with a particular body placement (for example, glasses are associated with the eyes), changes can still occur (the user might put her glasses in a shirt pocket or carry them in her bag).

Published by the IEEE CS ■ 1536-1268/14/$31.00 © 2014 IEEE

The influence that the body part on which the device is placed has on signals from motion sensors can be twofold: one, some activities are associated with specific body parts and sensors, such as activities related to subtle arm motions (washing hands), and two, even for activities that aren't strictly body-part specific, the motion sensor signals vary significantly from body location to body location (see Figure 1a).

A device's on-body location also impacts sound. As Figure 1b indicates, differences arise less from body damping and more from the absorption spectra of the clothing on which a device is placed.

Displacement within a Body Part

Depending on the specific piece of clothing, a device can move within a pocket or holster and thus might be at a slightly different location. In addition, a device can shift around if it is significantly smaller than the pocket or if the holster itself can shift (see Figure 2). The problem of "within body part" variations differs in three ways from the problem of inferring on which body part a device is placed:

- The placement variations are continuous rather than discrete.
- The variations are mostly small, usually not more than a few centimeters and at most 10 to 20 cm.
- The placement is likely to change more frequently as the device shifts around. Therefore, recognizing the within-body location is a difficult problem, making understanding how placement affects the sensor signal and defining features that are tolerant to it a better approach.

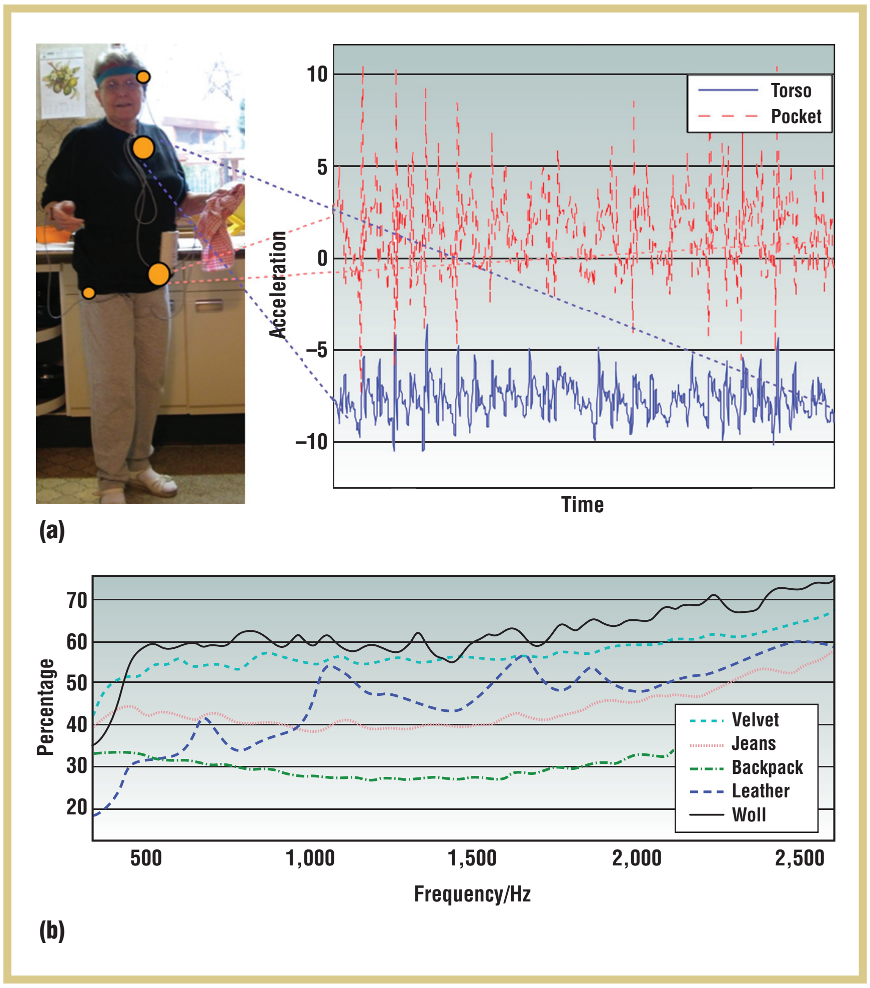

Figure 1. Impact of device placement on signals from different types of sensors. (a) Acceleration signals recorded from two different body locations while walking. (b) The sound absorption spectra of different types of clothing recorded with an iPhone microphone. The graph shows signal intensity as a percentage of the intensity recorded with a device placed on a table and not obstructed by clothing.

To investigate the impact of within-body displacement, we conducted experiments in two gym scenarios: one studying locomotion and the other arm exercises.[4] For each scenario, we used four Xsense inertial measurement unit (IMU) sensors uniformly and randomly distributed over the upper leg for the locomotion study and the upper arm for the arm exercises study. The locomotion exercises included eight classes: walking, running, running uphill, biking, rowing, walking up stairs, skiing, and cross-training. The arm exercises included using the lat/pectoral machine to do shoulder presses, upper back stretches, arm extensions/curls, pull downs, and chest presses. Two test subjects performed two circuits, with 20 repetitions of each exercise.

Using just accelerometer data from the same location for training (34 percent) and testing (66 percent), we achieved 100 percent recognition for the locomotion scenario and 96 percent for arm exercises. However, when testing and training on a different location, the performance dropped to 63 and 24 percent, respectively (see Table 1). Training the system on two different locations improved the recognition rate, but only slightly, to 61 and 31 percent.

Orientation

The final displacement issue is the device orientation with respect to the user's body. As an example, consider a mobile phone in a hip holster. In most cases, you can put the device into the holster in at least two ways: display

Figure 2. An example of "within body part" displacement. The placement of the device, although located on the arm, can shift up or down.

facing inward or display facing outward; you might also have the option of attaching the holster either vertically or horizontally, or be able to place the holster at various locations on the hip, which would determine its exact horizontal orientation.

In most cases, motion sensor-enabled devices are equipped with three-axis sensors, letting us use the signal's norm as a simple orientation-invariant signal. In this case, the interesting question is not "What influence does rotation have on motion signals?" but rather "How much information is lost when discarding orientation-sensitive features?" Depending on the antenna/microphone, the device's orientation on the body can dampen the signals for GPS, Wi-Fi, and sound.

TABLE 1
Classification results for displaced sensors in a simple gym experiment.

Modality	Trained and tested on same location	Trained on one location and tested on another	Trained on two locations prior to testing
Acceleration	100%	33%	35%
Gyroscope	65%	43%	44%
Cut off*	—	42%	47%
Combined†	—	78%	85%

*Just the acceleration features, but ignoring parts where displacement was detected
†Mixes acceleration features for translation and gyroscope from rotation-dominated body parts

Signal damping. Fixed rules for the influence of orientation on signal damping are difficult to formulate because the effect strongly depends on device and clothing configuration. Considering sound, signal dampening can increase if the microphone is oriented toward the user rather than the environment (see Figure 3).

Loss of orientation information. To show what kind of information is lost when discarding orientation, we used the standard locomotion recognition problem: distinguishing between walking upstairs, walking downstairs, running, or walking on a level surface. In Figure 4, we use Euclidean distance as a measure of time-series similarity to compare how well those classes are separated by the norm and any orientation-sensitive signals from an accelerometer signal's three individual axes. Even for simple classes (which can actually be recognized reasonably well using the norm), the orientation-sensitive features contain a lot more information.

On its own, device orientation is particularly relevant for systems containing a magnetic field sensor, as is increasingly the case with modern smartphones. Knowing the orientation of the magnetic sensor axis with respect to the user's body means that the sensor can be used to determine the direction in which the user is facing. In turn, this can be used to infer the user's focus of attention (such as facing a specific shelf in a store, thereby indicating interest in a certain product). More details appear elsewhere.[5]

Mitigating Placement Variations

We've discussed the types of sensor displacement and their effect on varying recognition modalities, so now we build on this conceptual framework to present a comprehensive set of methods for dealing with such effects. We don't provide a classical, exhaustive survey: there's a lot of current interest in this topic but not much work in the field yet. Our goal is to help activity recognition researchers make their systems more robust with respect to sensor displacement. For practical reasons, we provide examples from our own work to illustrate key concepts, but we take care to include relevant related approaches by others as well.

Addressing On-Body Variations

To mitigate the influences of on-body placement, we can apply either location-independent features or body part placement recognition.

Location-independent features and classifiers. For activity recognition applications, two strategies are commonly used to achieve location independence.

First, for some recognition tasks, we can compute features that are independent of the body part on which the device is placed. For detecting walking,

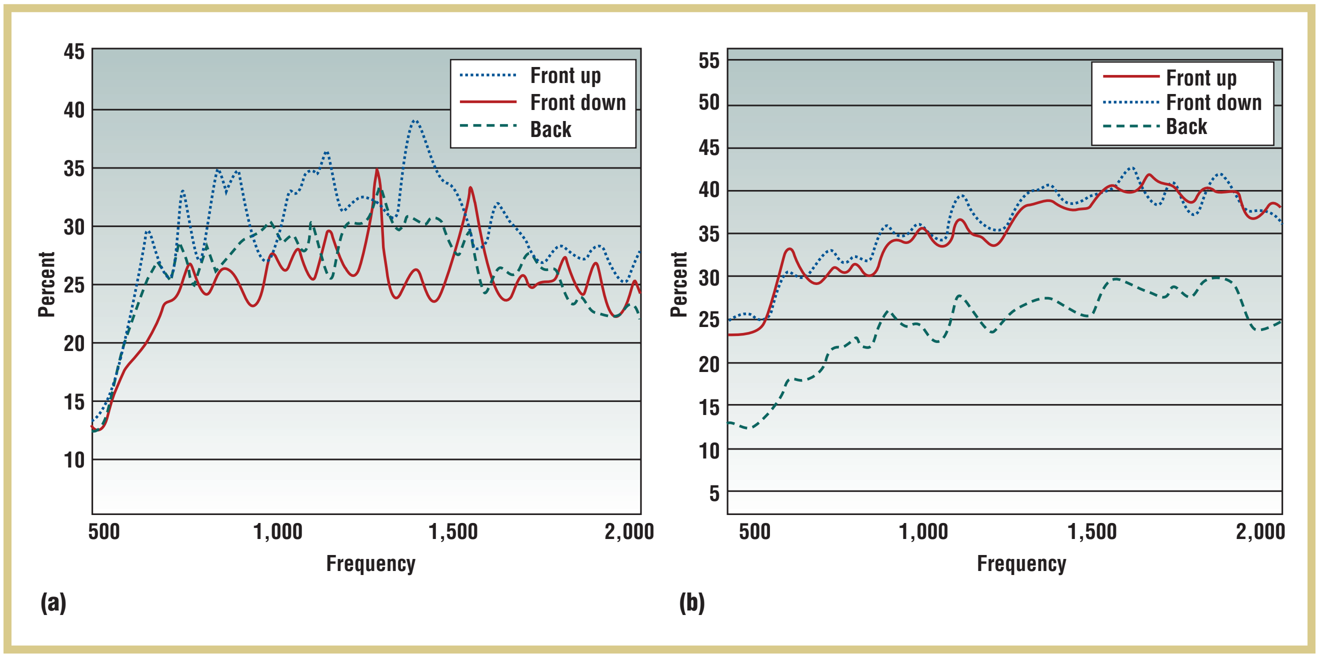

Figure 3. The influence of orientation on the audio signal absorption spectrum for the (a) Nokia 81 and (b) iPhone 3gs. Depending on where the microphone is located on the device, a different orientation will cause it to be more or less obstructed by the body or clothing.

for example, strong peaks between 1 and 4 Hz in the signal's frequency spectrum are present in all body locations. Our previous work shows how to detect walking independent of body location with up to 95 percent accuracy using the acceleration signal norm.[6] Other work focuses on how to use sensor fusion methods to achieve robust recognition.[7,8] As expected, these methods require significantly larger, more representative datasets for training compared to the non-robust alternative.

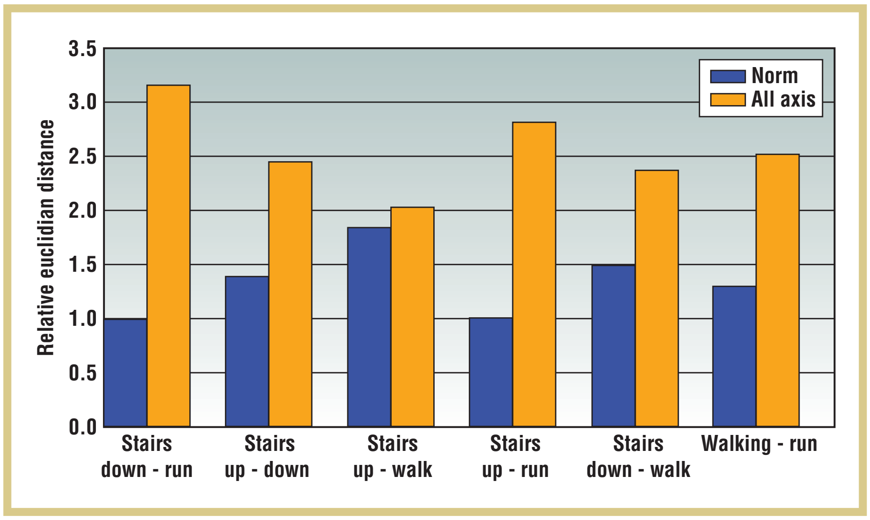

Figure 4. The value of the orientation information for motion sensor-based activity recognition, depicting the difference between recordings for different pairs of modes of locomotion as expressed by the vector norm (an orientation-invariant feature) of the acceleration on the upper leg and the corresponding three vector components. Note that the vector components can only be used if the device orientation is known.

Second, classifiers can be trained on a combination of different locations. Thus, the classifiers automatically select a feature set that's as independent as possible of the sensor location.[9] One study[7] examined such an approach for a multimodal sensor system that included different sensors (motion, sound, light), three different body locations (shoulder, wrist, hip), and eight everyday activities (different modes of locomotion). The classifiers trained and tested on different locations had up to 20 percent worse recall and precision than the ones trained and tested on the same location. A classifier trained on all three locations was only a few percent worse than the location-specific classifiers.

We investigated a similar approach with respect to sound classification in different pockets.[10] In the experimental setup, we considered sound from

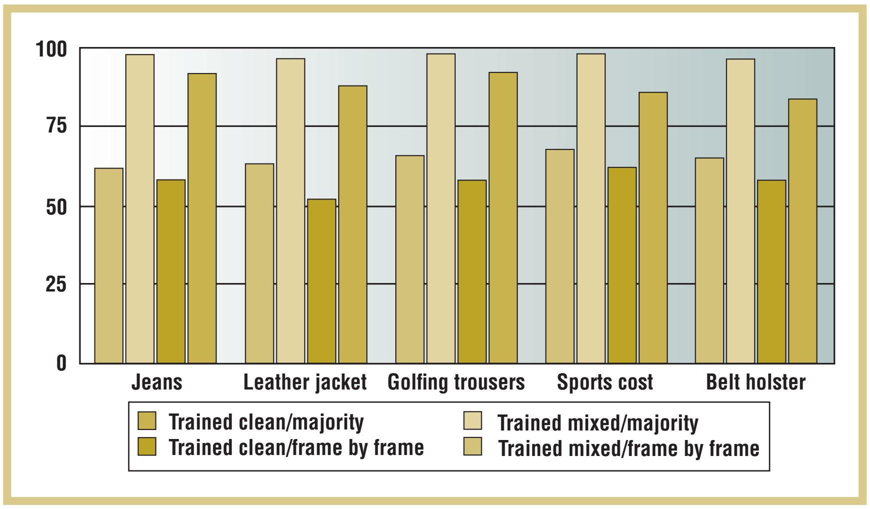

Figure 5. Classification results for audio from nine household and office appliances, showing the effect of sound dampening on accuracy. "Trained clean" refers to training with the phone being held in one's hand without absorption through clothing. "Trained mixed" refers to training on data from different pockets. "Majority" implies a majority decision over all frames in a recording instance.

nine household and office appliances recorded with an iPhone placed in two different trouser pockets, two different jacket pockets, a belt holster, and the user's hand. If we trained and tested the device at the same location, the recognition results were 99 percent (see Figure 5). If we trained the system with the device in hand—that is, without absorption through clothing (designated "clean" in Figure 5)—and tested it on different pockets, then the performance went down to around 60 percent on a frame-by-frame basis, and around 70 percent for the majority decision over all frames in a recording instance. By training on data recorded from different types of pockets, the system learns to deal with absorption through clothing, and performance can be improved to between 80 and 90 percent on a frame-by-frame basis and 100 percent for the majority decision.

Body placement recognition. Clearly, placement-invariant features and mixed training works only if there is sufficient similarity between signals from different body locations. However, if the signals differ sufficiently, we can automatically detect device placement. In previous work, we showed that such recognition is possible using the acceleration signal alone.[6] We were able to detect different on-body placement for torso, hand, leg, and head by relying on the fact that these parts all move with different frequency and along different distinct axes.

Leonard Grokop and his colleagues pursued a different approach to location recognition by using a fusion of acceleration, proximity, and light sensors to jointly recognize simple activities and on-body locations.[11] Thus, walking with the sensor in a pocket is a separate class from walking and having the sensor in the hand, reflecting the fact that both are separate events in signal space. A disadvantage of this approach is the explosion of the state space.

Related research focuses on phone placement sensing. Jason Wiese and his colleagues extended the work to detecting a phone's on-body location under real-life scenarios by utilizing accelerometer data and a capacitive sensor prototype. They also explored where people usually wear phones to adjust priors on different locations.[3] Emiliano Miluzzo and his colleagues presented the Discovery system, which can detect whether a phone is in a pocket by utilizing audio with a multiround classifier method.[12]

Addressing Within-Body Displacements

Features tolerant to within-body displacement of gyroscope and acceleration sensors[4] can be derived by assuming a rigid body idealization. Neglecting deformation assumes the distance between any two given points to be constant regardless of the external forces exerted on the object in question.

As Figure 6 shows, a rigid body can move in two ways: through translation and through rotation:

- During translation, all points in a rigid body move at exactly the same speed and acceleration. Thus, signals from accelerometers are location invariant. Because a translation per definition contains no rotational component, gyroscopes produce no signal.
- During rotation, each point in a rigid body moves with the same angular-velocity and acceleration. The gyroscope signal is the same for all points within the body, but the acceleration signal is placement dependent.
- In addition to the dynamic components resulting from motion, accelerometer signals contain a component related to gravity that's effectively identical at all places on an object and is thus invariant to displacement within body parts. It does, however, depend on sensor orientation.

From these facts, we can draw three conclusions:

- When motions are dominated either by translations or changes in orientation with respect to gravity, we can use acceleration features.
- When motions are dominated by rotations, we should avoid acceleration features; gyroscopes provide information that's invariant to body part displacement.
- When motions contain an equal amount of rotation and translation,

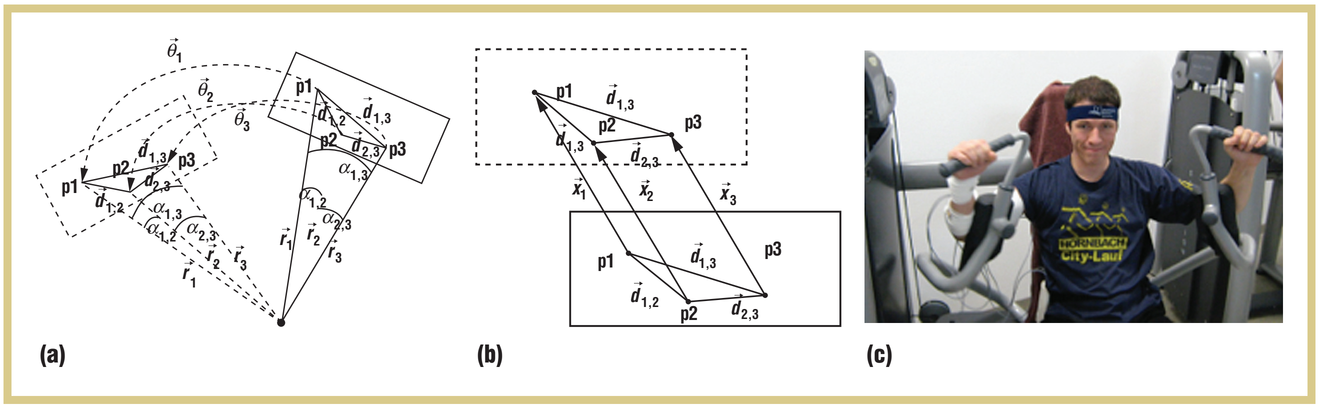

Figure 6. Rigid body approximation: (a) rotational and (b) translational motion of different points within a rigid body. It can be seen that a translational motion moves all points within a rigid body by the same distance, which means that they all have the same speed and acceleration at all times. By contrast, during rotational motion, (a) different points have different degree of displacement, which leads to different acceleration values (acceleration signal is location dependent).

> we should still pick the gyro features—although we might lose some information, we'll retain the shift invariance.

Unfortunately, arbitrary combinations of rotations and translations with a varying vertical component are still likely to occur. Consequently, static optimization will have only a limited effect on a typical recognition application's displacement invariance. To handle this, we need to choose the correct features dynamically, depending on the kind of motion being performed.

In previous work, we showed how to accomplish this kind of dynamic selection,[4] using the ratio of the acceleration vector norm to the angular velocity vector norm to determine if a signal frame is dominated by translation or rotation. As described earlier, translations generate no gyro signal, so the more the ratio is tilted toward the acceleration norm, the more translation-dominated the signal frame.

When applying this method to the gym experiment, we picked the decision boundary to be 150 degrees per second. Above this boundary, we used gyro features and set the accelerometer features to zero; below it, we discarded the gyro features and used the accelerometer features only. We made this decision for every data point at a 100-Hz sampling rate. As Table 1 shows, significant improvement can be achieved.

The standard practice for dealing with displacement is to take robust aggregate features. However, in general, such features only work for very simple recognition tasks. For more complex tasks, they lead to a significant degradation in recognition performance and require larger training datasets.[7,8] Navid Amini and his colleagues used robust rotation features to compensate for displacement while detecting a specific medical device's on-body location.[13] Kristof Van Laerhoven and colleagues trained recognition models to be adaptive to small placement issues. The classes were trained as user-dependent, under the assumption that a particular displacement happens more often with a specific user, one drawback being that this training required direct user feedback.[9]

K. Forster and his colleagues and H. Sagha and his colleagues regarded displacement as a continuous tracking problem.[14,15] The Sagha study also detected anomalies in the recorded data that the inference algorithm should disregard, but for their algorithm to work, it had to know the device's initial exact position on the body.

Addressing Orientation Variations

Estimating a device's orientation with respect to a user's body involves two distinct subproblems: vertical orientation (the angle with respect to the gravity vector) and orientation in the horizontal plane.

As first proposed by David Mizell and his colleagues, vertical orientation can be estimated when the object experiences no change in motion speed.[16] In this case, the only acceleration registered by the sensor is the Earth's gravity (9.81 m/s^2), and the direction of the measured acceleration vector defines the vertical plane. To identify signal segments with the above characteristics, the norm of the measured acceleration vector is used together with its variance. When variance of all axes tends toward 0 and the norm vector approaches 9.81 m/s^2, the signal is very likely to be dominated by the vertical orientation component. In theory, this might not necessarily be true: the object could be freely falling (and therefore, lack a gravity component) while experiencing a constant 9.81 m/s^2 along an arbitrary direction.

To derive the horizontal orientation of a body-worn accelerometer from a signal segment recorded while a user is walking, we have presented a method[25] based on the fact that, while walking, most variations in the acceleration signal's horizontal component are likely to be parallel to the direction of motion. Thus, we project the signal into the plane perpendicular to the vertical gravity vector (that is, the horizontal plane)

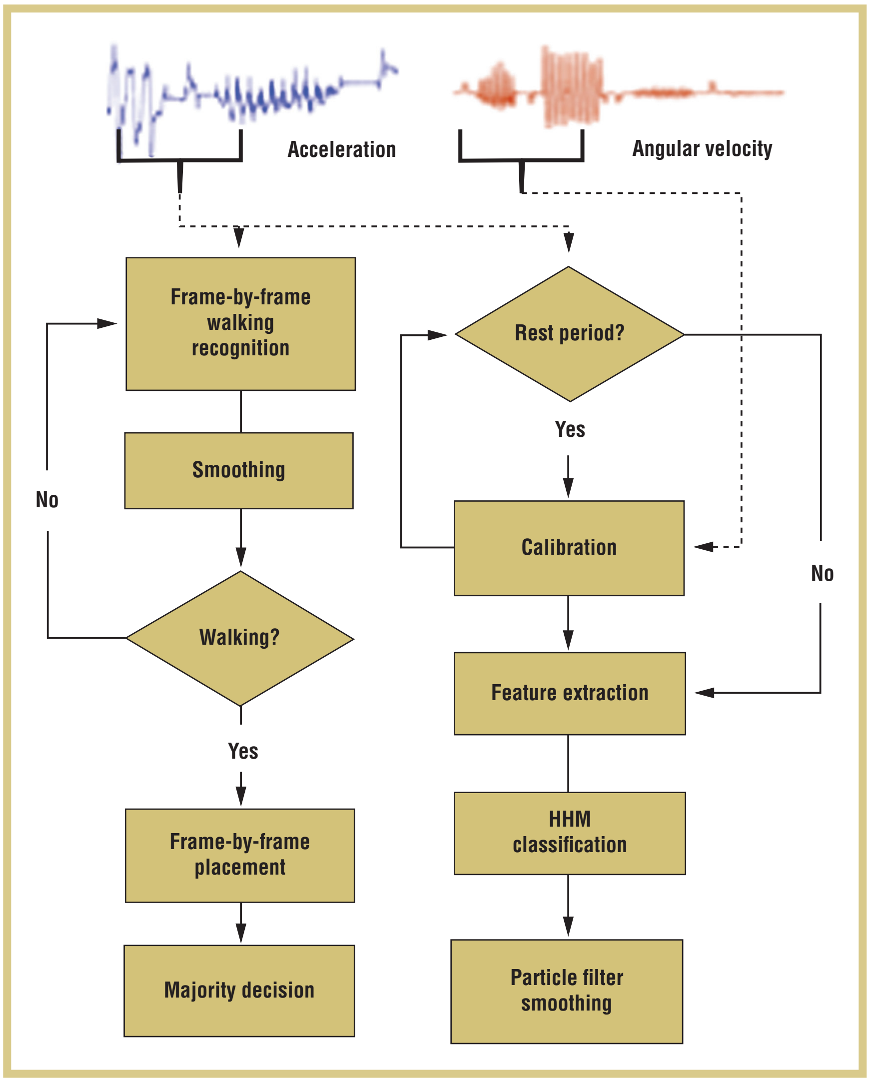

Figure 7. Performance of particle filtering for the respective datasets over time. It can be seen that the system needs approximately two minutes to reach its peak accuracy.

and apply principal component analysis to determine the direction in which the variation is greatest. To distinguish between front and back, we consider the integral of the signal over time.

Other Approaches

Most other methods for dealing with orientation changes focus on using orientation-robust features. Forster and colleagues describe how to track orientation and displacement changes over time, but they assume known sensor positions at the beginning.[14] Ricardo Chavarriaga and his colleagues[17] propose a similar yet more elaborate approach, an unsupervised adaption method that can compensate for small sensor displacements and orientation changes over time (this approach can't handle large changes, such as a user taking out his mobile phone and putting it back in a different orientation).

Surapa Thiemjarus and his colleagues defined device orientation as a classification problem,[18] implying that each different device orientation required its own training data. Oresti Banos and his colleagues presented a weighted sensor fusion-based approach,[19] while Ming Jiang and colleagues calculated a transformation matrix by using Gram-Schmidt orthonormalization to eliminate the sensor's orientation error and then employing a low-pass filter with a cut-off frequency of 10 Hz to eliminate the main effect of the sensor's misplacement.[20]

Additional dedicated work is in the area of modeling orientation variations in garment-integrated sensors.[21,22] The research here focuses on exploiting garment properties to estimate orientation changes and displacement.

Detecting On-Body Placement

Up to now, we've merely sketched possible solutions and methods (space doesn't permit a more detailed discussion). To provide some deeper insight, we pick a specific example and detail it here—specifically, recognition of body part location—and by extending our own prior work,[6] we describe how a recognition system looks in detail.

Different body parts vary in degrees of freedom and move in different patterns. We evaluated more than 35 features and found that six best capture those differences. The first three are standard deviation, zero crossings, and mean of the norm of the acceleration vector minus the gravitational pull g_0:

$$\left|\sqrt{x^2+y^2+z^2}-g_0\right|.$$

The fourth is the sum of the norm of the differences in variance for the normalized axes a_1, a_2, a_3 divided by the variance of the vector norm

$$\frac{1/2\sum_{i=1}^{n}\sum_{j=1,j<i}^{n}\left|var(a_i)-var(a_j)\right|}{var(norm)}.$$

The fifth is the number of peaks in the absolute value of the three axes derived using hill climbing with a threshold; the sixth is the median of these peaks.

We computed these features over a 2.5-sec jumping window (overlapping 1.25 sec); we then applied another six-minute window on top of the already windowed features and fed them into a continuous hidden Markov model (HMM) with five hidden states.

We evaluated the system on five datasets:

- *Household work*. We gathered six hours of real-life activities from a 70-year-old housewife, a 50-year-old female office worker, and 28-year-old male student.
- *Opportunity*. The largest dataset recorded as part of the Opportunity EU project includes activities from everyday living such as making a sandwich, pouring coffee, and eating. Two users repeated the setup five times.[23]
- *Drink and work*. The drink and work dataset contains mostly sitting activities, such as working on a computer and consuming food and drinks. In total, we studied six subjects with four repetitions each; one experimental run took around 30 to 40 minutes.
- *Bicycle repair*. The last experimental setup included repair activities on a bike (attaching a tire, tightening screws, and so on) with six test subjects.[24]
- *Office Work*. A set of simple office activates such as typing, photocopying, or participating in a meeting.

The maximum accuracy achieved for the six-minute windows was 82 percent, with roughly 80 percent at five minutes, as depicted in Figure 7. Because we recorded unconstrained real-life activities for our experiments, the dataset contained many segments that had no significant motions characteristic of body parts that we could use for classification, which explains the relatively low recognition rate. If we took just walking segments into account, for example, the recognition rates improve to 94 percent.[6]

To improve our on-body placement recognition work, we tried several filtering approaches. As a majority window is too crude to filter out uncharacteristic movements, we applied a sequential Monte Carlo method, also called a particle filter; the basic method for the HMM recognition remained the same. When we applied the particle filtering, we could reduce the sliding window for each individual HMM to 45 seconds. We input these 45-second "sliding window" HMM classifications as observations into a particle filter. To our knowledge, particle filtering hasn't been applied to this type of activity-sensing problem. Figure 8 gives an overview of all methods used.

To test our filter design, we randomly (uniformly) picked 100 10-minute segments from the datasets (there were duplicates), did feature extraction and HMM classification, and then fed them into the particle filter. Afterward, we evaluated what percent of the 100 filtered placements we detected correctly. Figure 7 shows the results for the different datasets. As can be seen, the office work scenario performed best (using gyroscope and accelerometer), and the opportunity dataset was second (gyroscope and accelerometer). House work started off quickly (accel + gyro), but didn't achieve the average 90 percent accuracy and had the most variance from one filtered classification to the other. For all datasets, we reached a recognition rate of more than 90 percent after 2 to 4 minutes.

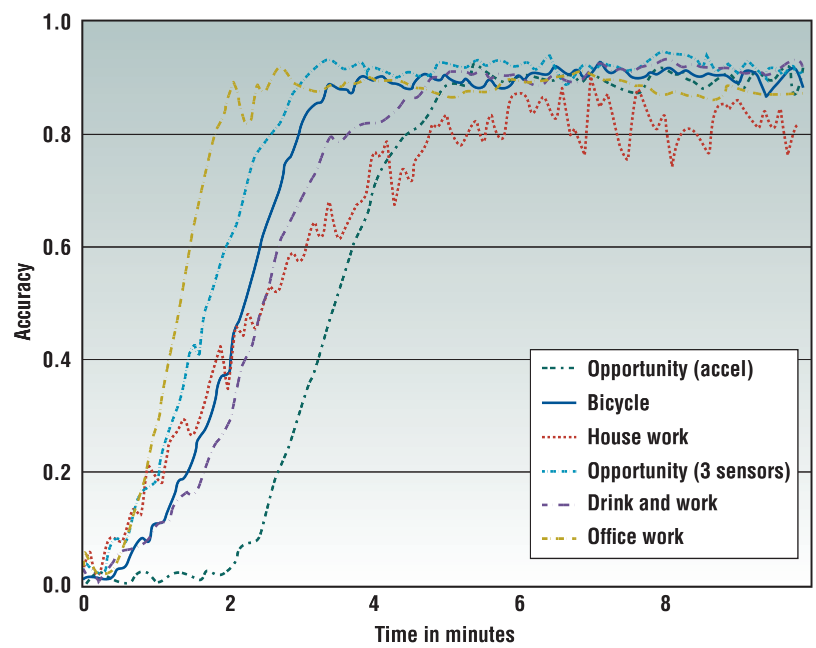

Figure 8. Method overview for the on-body placement recognition based on walking (left) and unconstrained using HMMs/particle filtering (right).

TABLE 2
On-body recognition rates (%).

	Head	Wrist	Torso	Front pocket	Back pocket
Head	74.0	1.0	12.5	0	12.5
Wrist	0	87.2	6.4	2.6	3.8
Torso	7.1	5.9	84.7	0	2.4
Front pocket	0	2.1	14.9	37.2	45.7
Back pocket	0	2.0	11.9	2.0	84.2

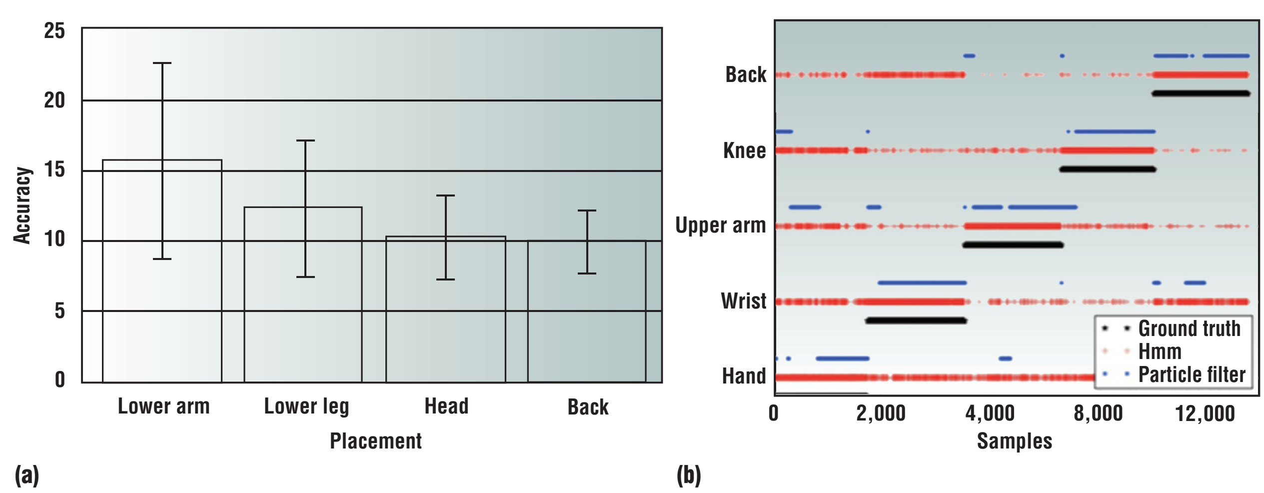

Figure 9. On-body location recognition using HMMs and particle filter smoothing: (a) Mean and standard deviation of the time it takes before the particle filter recognizes the correct placement using 100 segments between 10 and 20 minutes, uniformly randomized for the opportunity dataset. (b) Scatter plot depicting HMM performance with and without particle filter smoothing; 20-minute segments taken from the opportunity dataset accelerometer only, with 45-sec. HMM classification (around 59 percent correct) and achieving a 78 percent particle filter with 40 particles.

The recognition for body parts with less movement didn't perform that well (head and torso), as can be seen from the confusion matrices in Table 2 and Figure 9a. The classification was worse for the HMMs alone, but it can be significantly improved by applying the particle filter (Figure 9b).

We discussed many issues here with respect to smartphones and mobile appliances, but our points are also relevant for other types of wearables—in particular, the body part shifts and rotations that are well-known problems for smart textiles, which are never firmly fixed to the body (sleeves may be rolled up or sweatshirts carried around the hip or in a backpack). ■

REFERENCES

1. T. Choudhury et al., "The Mobile Sensing Platform: An Embedded System for Capturing and Recognizing Human Activities," *IEEE Pervasive Computing*, vol. 7, no. 2, 2008, pp. 32–41.
2. K. Forster et al., "Evolving Discriminative Features Robust to Sensor Displacement for Activity Recognition in Body Area Sensor Networks," *Proc. 5th Int'l Conf. Intelligent Sensors, Sensor Networks and Information Processing* (ISSNIP), 2009, pp. 43–48.
3. J. Wiese, T. Scott Saponas, and A.J. Brush, "Phoneprioception: Enabling Mobile Phones to Infer Where They Are Kept," *Proc. SIGCHI* (CHI 13), 2013, pp. 2157–2166.
4. K. Kunze and P. Lukowicz, "Dealing with Sensor Displacement in Motion-Based On-Body Activity Recognition Systems," *Proc. UBICOMP*, 2008, pp. 20–29.
5. B. Steinhoff and U. Schiele, "Dead Reckoning from the Pocket: An Experimental Study," *Proc. PerCom*, 2010, pp. 162–170.
6. K. Kunze et al., "Where Am I: Recognizing On-Body Positions of Wearable Sensors," *Location and Context Awareness*, Lecture Notes in Computer Science, vol. 3479, 2005, pp. 264–275.
7. J. Lester, T. Choudhury, and G. Borriello, "A Practical Approach to Recognizing Physical Activities," *Proc. 4th Int'l Conference Pervasive Computing*, 2006, pp. 45–52.
8. A. Krause et al., "Unsupervised, Dynamic Identification of Physiological and Activity Context in Wearable Computing," *Proc. Int'l Symp. Wearable Computers*, 2003, pp. 36–42.
9. K. Van-Laerhoven and O. Cakmakci, "What Shall We Teach Our Pants?" *Proc. Int'l Symp. Wearable Computers*, 2000, pp. 77–83.
10. T. Franke et al., "Can a Mobile Phone in a Pocket Reliably Recognize Ambient Sounds?," *Proc. Int'l Symp. Wearable Computers*, 2009, pp. 161–162.
11. L.H. Grokop et al., "Activity and Device Position Recognition in Mobile Devices," *Proc. UBICOMP*, 2011, pp. 591–592.
12. E. Miluzzo et al., "Pocket, Bag, Hand, etc.: Automatically Detecting Phone Context through Discovery," *Proc. PhoneSense*, 2010, pp. 21–25.
13. N. Amini et al., "Accelerometer-Based On-Body Sensor Localization for Health and Medical Monitoring Applications," *Pervasive and Mobile Computing*, vol. 7, no. 6, 2011, pp. 746–760.
14. K. Forster, D. Roggen, and G. Troster, "Unsupervised Classifier Self-Calibration through Repeated Context Occurrences: Is There Robustness against Sensor Displacement to Gain?" *Proc. 10th Int'l Symp. Wearable Computers*, 2009, pp. 77–84.
15. H. Sagha et al., "Online Anomaly Detection and Resilience in Classifier Ensembles," *Pattern Recognition Letters*, 2013.
16. D. Mizell, "Using Gravity to Estimate Accelerometer Orientation," *Proc. 7th Int'l Symp. Wearable Computers*, 2005, pp. 252–253.

17. R. Chavarriaga, H. Bayati, and J. del Millan, "Unsupervised Adaptation for Acceleration-Based Activity Recognition: Robustness to Sensor Displacement and Rotation," *Personal and Ubiquitous Computing*, vol. 17, no. 3, 2013, pp. 479–490.

18. S. Thiemjarus, "A Device-Orientation Independent Method for Activity Recognition," *Proc. Body Sensor Networks*, 2010, pp. 19–23.

19. O. Banos et al., "On the Use of Sensor Fusion to Reduce the Impact of Rotational and Additive Noise in Human Activity Recognition," *Sensors J.*, vol. 12, no. 6, 2012, pp. 839–854.

20. M. Jiang et al., "A Method to Deal with Installation Errors of Wearable Accelerometers for Human Activity Recognition," *Physiological Measurement*, vol. 32, no. 3, 2011, p. 347.

21. H. Holger Harms, O. Amft, and G Troster, "Modeling and Simulation of Sensor Orientation Errors in Garments," *Proc. BSN*, 2009, pp. 20–26.

22. G. Gioberto and L.E. Dunne, "Garment Positioning and Drift in Garment-Integrated Wearable Sensing," *Proc. Int'l Symp. Wearable Computers*, 2012, pp. 64–71.

the AUTHORS

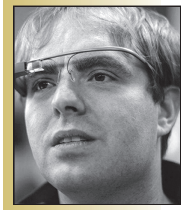

Kai Kunze is a research assistant professor in the Intelligent Media Processing Group at Osaka Prefecture University. His research includes the exploration of knowledge acquisition tasks, combining computer vision and other embedded sensing approaches with a focus on reading activities and learning behavior. Kunze received a PhD from University Passau, Germany. Contact him at kai.kunze@gmail.com.

Paul Lukowicz is a professor of artificial intelligence at the Technical University of Kaiserslautern Germany, and scientific director for embedded intelligence at the German Research Center for Artificial Intelligence (DFKI). His research interests are in wearable systems, ubiquitous and crowdsourced sensing, activity and context recognition, sociotechnical systems, and self-organization. Lukowicz received a PhD in computer science from the University of Karlsruhe. He is a member of the editorial board of the *Hindawi Advances in Human Computer Interaction* and *IEEE Journal of Biomedical and Health Informatics*. Contact him at paul.lukowicz@dfki.de.

23. D. Roggen et al., "Opportunity: Towards Opportunistic Activity and Context Recognition Systems," *Proc. Int'l Symp. World of Wireless, Mobile and Multimedia Networks* (WowMoM), 2009, pp. 1–6.

24. G. Ogris et al., "Continuous Activity Recognition in a Maintenance Scenario: Combining Motion Sensors and Ultrasonic Hands Tracking," *Pattern Analysis Applications J.*, vol. 15, no. 1, 2012, pp. 87–111.

25. K. Kunze et al., "Which Way Am I Facing: Inferring Horizontal Device Orientation from an Accelerometer Signal," *Proc. Int'l Symp. Wearable Computers* (ISWC 09), 2009, pp. 149–150.