Christian Debes, Andreas Merentitis, Sergey Sukhanov,
Maria Niessen, Nicolaos Frangiadakis, and Alexander Bauer

# Monitoring Activities of Daily Living in Smart Homes

*Understanding human behavior*

Monitoring the activities of daily living (ADLs) and detection of deviations from previous patterns is crucial to assessing the ability of an elderly person to live independently in their community and in early detection of upcoming critical situations. "Aging in place" for an elderly person is one key element in ambient assisted living (AAL) technologies.

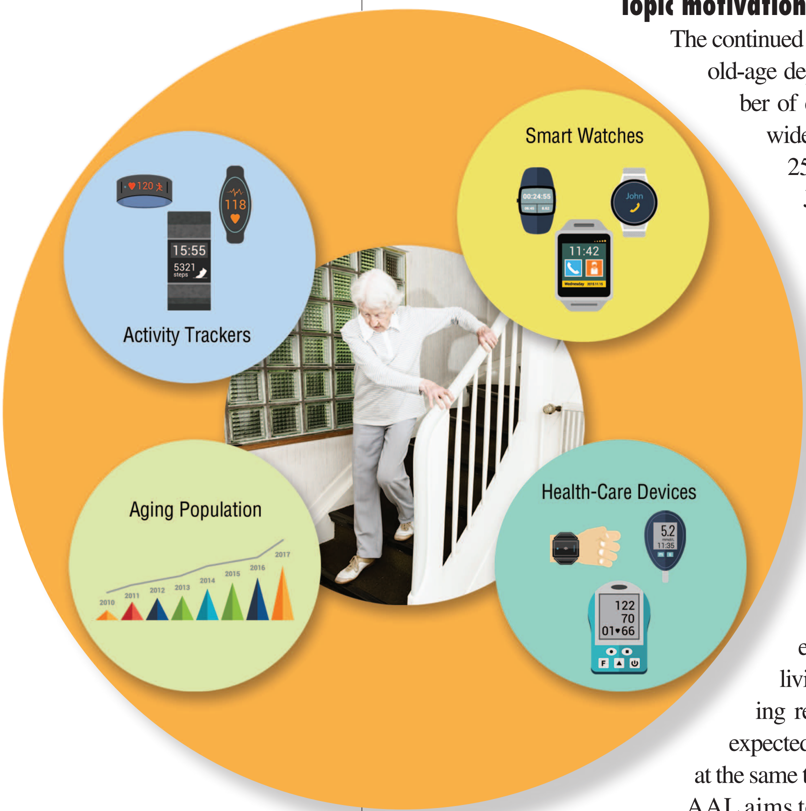




## Topic motivation and significance

The continued increase in longevity will yield a steep rise in the old-age dependency ratio, defined as the ratio of the number of elderly people to those of working age. Worldwide, this ratio is expected to double from 11.7% to 25.4% in the next 35 years, with countries like Japan, Germany, Italy, Spain, and Poland exceeding a 50% ratio [1]. In alignment to this development, the number of people aged 80 and over is going to triple in the next 35 years, going from 22 million to 61 million in the European Union [2] with similar developments in the United States [3].

This trend leads to several sociological and economical challenges. On the one hand, several studies show that, at 90%, the vast majority of elderly people have the desire to live as long as possible, independently, in their own home [4]. On the other hand, there is the desire of families and health insurers to have cost-effective alternatives to assisted living and nursing homes. The costs of maintaining retirement living standards due to longevity are expected to roughly double in the next 35 years, while, at the same time, a shortage of caregivers is expected [5], [6].

AAL aims to deal with some of the challenges that develop with longevity. It serves as a framework of solutions ranging from medication reminder tools to fall detection systems and communication tools. The technology used in these solutions is based on ambient intelligence, a paradigm within information technology that aims to aid people in their everyday lives by learning and adaptively responding to their behavior by integrating technology in their environment. As such, it can also assist elderly people to age in place while still having sufficient security standards in case of emergency [7]–[9].

*Digital Object Identifier 10.1109/MSP.2015.2503881*
*Date of publication: 7 March 2016*

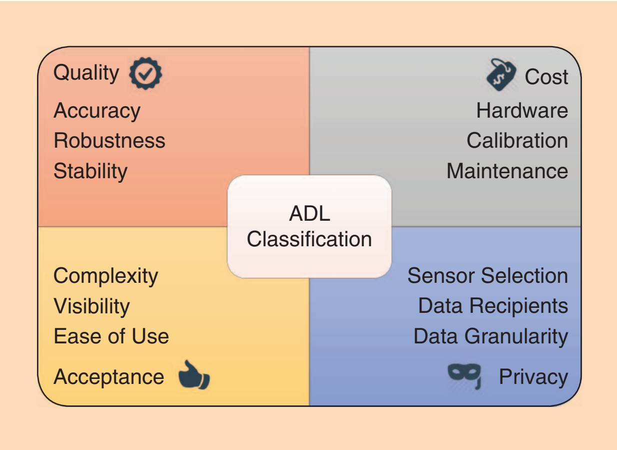


**FIGURE 1.** The challenges in ADL classification.

One important concept in AAL is the monitoring of ADLs [10]. The concept of ADLs is commonly used in health care, summarizing activities and daily routines, on which the functional status of a person is based, and, ultimately, on which the ability of a person to live independently in a community is assessed. These include six basic ADLs (bathing/showering, dressing, feeding, functional mobility, personal hygiene, and continence) meant to assess physical self-maintenance and a larger number of instrumental ADLs, such as food preparation and housework.

> **A primary goal of AAL is to assess the self-maintenance of elderly people still living at home.**

Because the manual assessment of the ADLs of a person is not feasible in a real-life situation, automatic classification and monitoring of ADLs using sensors deployed in households is a crucial technology for AAL. ADL monitoring can allow for early detection of diseases such as Alzheimer's [11], [12] and dementia [13], [14] and can generally reveal a decrease in the ability of a person living independently. ADL monitoring yields several technical and nontechnical issues that need to be addressed. On the technical side, the choice and setup of sensors deployed in households, as well as the signal processing and machine-learning algorithms to be considered for event detection and classification, are important. On the nontechnical side, ease of use and privacy are crucial [15], [16]. The most practically successful and useful system for ADL monitoring is thus one that requires little training or configuration effort and integrates seamlessly in a household. These considerations pose several challenges on the technology side, including:

- *Sensor selection.* Sensors have to be affordable, privacy preserving, and easy to install and configure, ruling out complicated sensors and microphones. This effects the achievable classification accuracy.
- *Household invariance.* Data and ground truth acquisition for each individual household is costly and laborious. ADL classifiers should provide reasonable performance on a variety of household configurations, with additional training data as optional input to boost accuracy.

Figure 1 summarizes some of the challenges that must be considered in a practically applicable ADL classification system.

## State of the art in sensor technology to assess ADLs

Reliable and accurate sensor data is crucial for ADL monitoring and classification tasks. Sensor effectiveness largely depends on the activity type to be recognized. In past works on ADL classification, various types of sensors were deployed in experiments leading to different architectures and performance of the overall systems [17]. Two main categories of sensors can be distinguished: wearable sensors and nonwearable sensors. Wearable sensors are usually attached to a person directly (e.g., bracelet sensors or cardio sensors) or to their clothes (e.g., an accelerometer or a step counter) to measure location, pulse rate, body temperature, blood pressure, and other vitally important metrics as well as motion characteristics. Nonwearable sensors are usually deployed in stationary locations of a house or a room and are able to detect a person and his movements and activities. Nonwearable sensors can specify the operational status of objects, measure water flow, room temperature, or door/cupboard openings/closings. While wearable sensors allow for higher localization accuracy and can detect body movements and vital health metrics [18], nonwearable sensors are considered less intrusive and do not require any interaction from the user's side. Wearable sensors also may have harsher power consumption requirements. However, in some cases, the wearable sensors might be part of or make use of devices the user already is familiar with and normally carries with them, such as a wristwatch or a cell phone.

### *Nonwearable sensors*

In Table 1, we summarize and categorize nonwearable sensors that were used for ADL monitoring and classification in previous work.

- Infrared (IR) sensors are the most often used nonwearable sensors in past projects and studies on ADL classification [19]–[24]. They are used to discover human presence in a room, detect motion in a specific area, or to locate a human within a house. In [25], a modified passive IR (PIR) sensor was used to detect stove and oven operation.
- Ultrasonic sensors are usually used for person detection and localization by measuring distances to objects. In [26]–[28], these sensors were deployed together with other sensors to monitor the behavior of a person and to identify ADLs. In other studies, ultrasonic sensors were used to get accurate pacing trajectories and then to find ones that were abnormal [29], [30].
- Photoelectric sensors are devices that detect a light source and output a signal when the light intensity is greater or less than the predefined threshold value. This type of sensor is not extensively used; however, in some projects, they

**Table 1. Nonwearable sensors used for ADL classification.**

| Sensor | Type of Measurement | Task | Usage Example |
|---|---|---|---|
| Passive/Active IR | Motion/Identification | Localization/Presence detection | Detection of person in kitchen |
| Ultrasonic | Motion/Identification | Localization/Presence detection | Detection of person in kitchen |
| Photoelectric | Motion/Identification | Localization/Presence detection | Detection of person in kitchen |
| Video/Thermal | Activity | Localization/Presence detection | Detection of person next to stove |
| Vibration | Vibration | Presence detection/Object usage | Detection of person in kitchen |
| Pressure | Pressure on object | Presence/Fall/Steps detection | Fall detection |
| Magnetic switches | Door/Cupboard opening/Closing | Objects usage/Presence detection | Cupboard opening |
| Radio-frequency identification (RFID) | Object information | Objects usage/Presence detection | TV usage |
| Audio | Activity | Objects usage/Presence detection | Shower usage |
| Wattmeter | Consumption information | Electrical objects usage | Water boiler usage |

are used as a presence detection sensor [31], [32] or as gait speed and direction measurer [33].

There are video-based approaches in which a camera is installed in a specific place of a house to detect person movements and/or other general activities. While performing well under laboratory conditions, this type of sensor is unable to provide the same performance in natural conditions because of noise and nonconstant lighting [34]. Moreover, a video-camera-based approach is considered to be strongly privacy-violating. To address privacy concerns, a low-resolution thermal sensor was recently proposed to be used instead of a traditional video camera [35], [36]. This sensor is able to provide almost the same activity information as a video camera while preserving the user's privacy. However, there are no studies that prove high operational performance of such types of sensors in real scenarios.

**A smart home is a normal living environment augmented with technology to improve the comfort or security of its residents.**

- Vibration sensors are usually deployed to detect a person falling [25], [37]. Vibration sensors can also be used in identifying interaction with various objects [38], flushing toilets, or detecting water flows [39], [40].
- Pressure sensors are used to detect the presence of a person, steps, and fall events. These sensors are usually deployed in the form of floor mats and smart tiles [25], [31]. In [41], pressure sensors were installed not only in floors but also in furniture to obtain object usage information during activities.
- Magnetic switches are usually used to report whether doors or cupboards are opened or closed. These sensors are also able to provide information on users accessing particular rooms and opening dressers, refrigerators, or trash cans. Details on installation and usage of magnetic switches and other types of door sensors can be found in [24], [31], [42]–[44].
- Audio sensors are usually used to detect sounds in-house and discriminate between different types of sounds. In [27], [45] microphones were installed to classify environmental sounds into classes such as speech, phone ringing, dish clanging, and TV/radio to extract events such as talking, a door closing, a person walking, a phone ringing, an object falling, and TV usage. In [46] an array of acoustic sensors was installed to detect a person falling.
- A Wattmeter and other sensors that measure electricity consumption of domestic appliances and light are often used in identifying ADLs. Today, this can be one of the major indicator of well-being of a subject [47]. In [48], electricity consumed by room lights and various appliances was used to record electrical activity and then to translate it into the probability of a particular ADL. In [49] domestic energy was monitored along with other sensors to find abnormalities and monitor the person's health and security status.

### *Wearable sensors*

In Table 2, we present a summary of wearable sensors that were used for activity recognition and ADL classification. Accelerometers are the most commonly used sensors for action, movement, and activity recognition [44], [50]. Attached to a specified human body, location accelerometers allow to differentiate between different types of motion (e.g., running, walking, sitting, scrubbing, etc.) [51], [52] or help to identity the posture of a person [53]. Often, accelerometers are also used to detect falls by measuring vibration or acceleration [54]. In some studies, accelerometers are used together with gyroscopes to obtain orientation information and better distinguish various types of motion and movement [55].

**Table 2. Wearable sensors used for ADL classification.**

| Sensor | Task | Usage Example |
|---|---|---|
| Accelerometer | Action recognition, types of motion, fall detection | Person running, person falling |
| Hand-worn sensors | Recognition of gestures, step counter | Person eating, person walking |
| Smartphone | Recognition of actions, movements, and types of motion | Person sleeping, person riding a bike |
| RFID | Recognition of actions of a person with objects | Cutlery usage, opening of cupboards, kitchen device usage |
| Vital monitoring sensors | Monitoring vital body parameters | High blood pressure detection, abnormal heart rate |

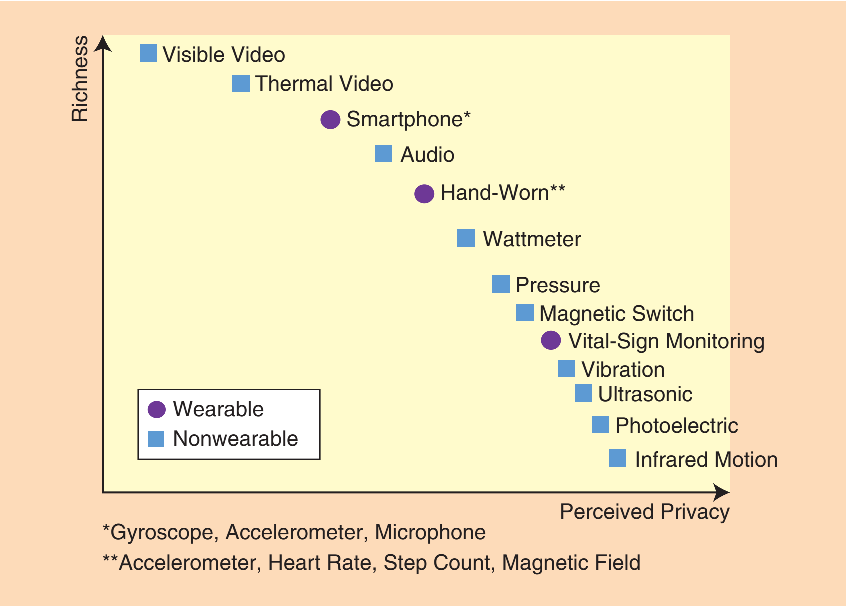


**FIGURE 2.** The richness of the sensors versus a user's perceived privacy.

Various types of hand-worn sensors are also considered in many activity recognition scenarios. These types of sensors include multifunctional wristwatches, magnetic sensors and other types of bracelets. Accelerometers are often integrated into wristwatches providing hand and arm gesture recognition capabilities [56]. In [57] an wrist-worn activity detector was used to perform sleep/awake activity classification. In [58] inertial sensors, accelerometers, and tilt switches were combined in a wrist-worn sensing unit to model users' rhythms and as a result recognize daily activities. Hand-worn magnetic sensors are able to distinguish between magnetic fields emitted by different electrical devices and recognize the activity of a user [59]. Emergency buttons are often integrated into a wristwatch and used to request help in case of an emergency situation [31]. Also, step counters are often integrated into wristwatches [57].

The modern smartphone offers a wealth of sensors, and can further be used as a communication platform [60]. Usually, accelerometers, gyroscopes, a global positioning system, a magnetometer and a microphone are incorporated into a modern smartphone device, providing all necessary information for movement, action and activity recognition including fall detection [61].

RFID tags are often used to detect the interaction of a person with an object and infer an ADL. In [62], RFID tags were deployed on various kitchen utensils such as bowls, cutlery, dishes, and jars to detect food preparation, eating, and drinking as well as on various cupboards, the TV, and furniture. Similar setups are considered in [63]–[65]. Matic et al. [66] focus specifically on monitoring dressing activity and detecting dressing failures. Often, RFID sensors are used in combination with other sensors such as accelerometers [66].

In addition to these sensors, there is also a large variety of sensors that monitor vital signs such as blood glucose, humidity and temperature, blood pressure, heart rat, pulse oximetry, $CO_2$ gas, electrocardiography, electroencephalography, electromyography, and electrooculography. These allow for the monitoring of a large set of human vital statistics and support activity recognition and ADL classification tasks [67], [68].

The ability of all of these sensors to provide rich information about people's lives and biometrics can raise severe privacy concerns. Figure 2 illustrates the richness of the sensors versus the perceived privacy of a person using these sensors. It is clear from the figure that sensors that provide rich information about a person are usually not perceived as privacy-preserving. For example, a video camera that allows for the recognition of almost any human activity in its field of view cannot be used in most rooms due to heavy privacy violations. In contrast, magnetic switches can be placed in every room without severely violating privacy, but they do not provide exhaustive information on every human activity. This can be partly improved by installing multiple instances of low-informative sensors (e.g., magnetic switches) so that richer insight into the user's activities can be achieved up to the level of the most informative sensors. However, this comes at the cost of increased installation efforts and complicates the deployment process.

## State of the art in ADL experimental setups

A primary goal of AAL is to assess the self-maintenance of elderly people still living at home. Therefore, many studies in ambient intelligence focus on automatically recognizing human activities that correspond to ADLs, such as bathing,

**Table 3. A summary of data sets collected in smart home environments, with their name and reference, whether wearable and/or nonwearable sensors were installed, approximately how long the (average) recording time was, and whether it was recorded in a (living) lab or a real home.**

| Data set | Institution | Sensor Types | Recording Duration | Lab/home |
|---|---|---|---|---|
| CASAS [70] | Washington State University | Wearable and nonwearable | Up to months | Lab and home |
| HIS [27] | Grenoble TIMC-IMAG Lab | Wearable and nonwearable | Hours | Lab |
| [23] | University of Virginia | Nonwearable | Weeks | Home |
| [24] | University of Amsterdam | Nonwearable | Weeks | Home |
| TigerPlace [71] | University of Missouri | Nonwearable | Year | Home |
| [65] | Intel Research Seattle | Wearable | Weeks | Home |
| [72] | Staffordshire University/Chiang Mai University | Wearable | Days | Lab |
| [63] | TU Darmstadt/Fraunhofer IGD | Wearable | Hours | Lab |

cooking, and eating, to be able to determine any changes in their patterns. The experimental setting in which human activity data can be collected is called a *smart home*. A smart home is a normal living environment augmented with technology to improve the comfort or security of its residents [69]. In the domain of AAL, sensors installed in the smart home can be used to monitor the behavior of people living in the home. For example, a team at Washington State University introduced the Center for Advanced Studies in Adaptive Systems (CASAS) Smart Home to test machine-learning techniques for human activity recognition [70].

Depending on the focus of a study, the experimental scenario and, consequently, the requirements on the smart home environment vary. The smart home can be a real home where sensors are installed, but it may also be a lab in which a smart home is built and where temporary residents can stay for a shorter or longer period of time. In addition, some studies use predefined scenarios to be able to systematically evaluate activity recognition algorithms, while others investigate patterns of normal behavior. Finally, the type of sensors that are installed vary, depending on the focus, e.g., energy efficiency or privacy considerations. Table 3 lists a selection of smart home data sets and properties of the experimental settings.

Related to the two types of sensors described in the "State of the art in sensor technology to assess ADLs" section—wearable and nonwearable—the experimental approaches can be separated in in-situ and ambient approaches. In the in-situ approach, the goal is to correctly identify particular activities, and this is often tested in a laboratory setting for a short period of time according to predefined scenarios. The types of sensors used are mostly low cost and low power, so that many can be installed. These include accelerometers [63], [73], both body-worn and attached to objects; RFIDs [74], [75], also both body-worn and attached to objects; and door contact sensors. Although wearable sensors allow experiments to include activities outside of a home, contrary to the ambient approach, most work in the in-situ approach and are applied indoors and in living labs. The advantage of using low-cost and low-power sensors is the possibility of running the installation for long uninterrupted periods. However, because these relatively simple sensors require a wide coverage, the initial set up requires more effort. Moreover, wearable sensors may not be easily accepted by elderly users.

The ambient approach is usually applied in experiments of longer duration in real-life settings, either in a smart home where participants live in an apartment (days or weeks, e.g., [70]) or in a real apartment (e.g., [71]). In the controlled environment of a smart home it is easier to gather detailed and balanced data and annotate them, for example, with cameras, making it suitable to gather data to test activity recognition algorithms. On the other hand, data recorded in real environments is more representative of normal behavior and therefore more suited to test algorithms for behavior modeling. For example, in [24], ambient sensors such as door contact sensors, motion sensors, and a float sensor in the toilet were used to recognize patterns of activities. This example was followed in [76] as part of the CASAS project to detect broad activities such as eating breakfast, sleeping, and wandering.

**Generative models estimate the joint probability distribution of observation samples, which can be used to predict the most likely class to which a new sample belongs.**

## ADL classification

The signal processing and machine-learning methods that are referenced in the literature on ADL classification span a broad range of techniques, from simple heuristics to more advanced machine-learning algorithms such as hidden Markov models (HMMs) and conditional random fields (CRFs). Most of the classical machine-learning algorithms such as support vector machines (SVMs) and random forests assume input data that is independent and identically distributed (IID). However, there are certain cases where the independence assumption of each data point does not hold. This is true, for example, in speech recognition (every syllable is dependent on the nearby ones) but also for human behavior modeling and recognition: What someone is doing at a specific point in time is not independent from what he was doing just before. The taxonomy of machine-learning algorithms that are used for structured learning when the IID assumption does not hold is presented in Figure 3.

Two broad categories in machine learning are generative and discriminative models, where the former is modeling the joint probability distribution of the samples and the labels and the latter is modeling the conditional probability of the labels given the samples. The standard HMM is a typical algorithm of the first category, with several of its extensions also falling into the same group. In the discriminative group, the basic models are CRFs and their extensions [for example, latent-dynamic CRFs (LDCRFs) and semi-Markov CRFs (SMCRFs)] as well as certain types of artificial neural networks (ANNs), with the most prominent ones being the recurrent neural networks (RNNs).

Finally, a multitude of hybrid methods, aiming to combine the advantages of discriminative and generative models, are also available. These include, for example, approaches relying on kernel metric distances such as the Fisher kernel and various combinations of HMMs with discriminant models such as random forests and ANNs.

While most work in ADL classification is performed using one of the aforementioned machine-learning techniques, heuristic methods also were successfully applied. Short-term activities and data sets with sufficiently redundant sensor setups (to suppress false alarms) are especially suitable for heuristic methods. One successfully applied heuristic is the circadian activity rhythms [23], [77], which describe the measurement of home rhythmic behavioral activity as the resident engages in the habitat. In some cases, these simple heuristics are either fused together or used as features for a second-level machine-learning algorithm. For example, in [78], simple heuristics measures like means and variances are used as features for neural network models, while the outcomes of the neural networks are fused under an HMM.

**Because HMMs are suitable to model sequential data, it is a popular classification method in activity recognition.**

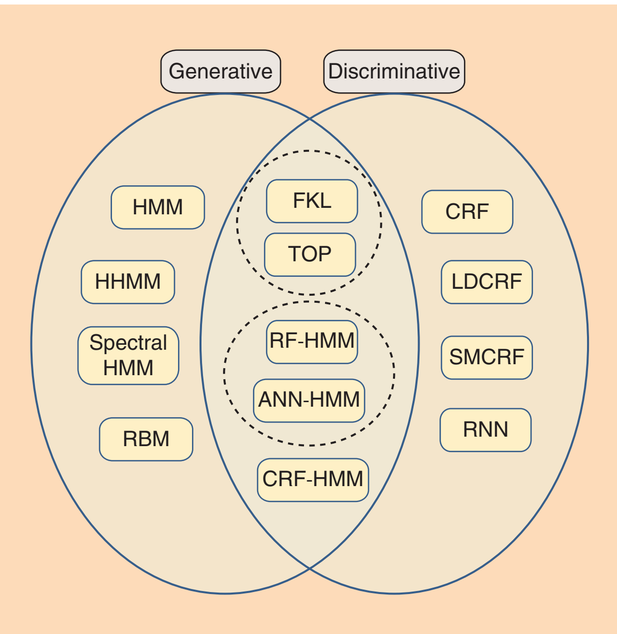


**FIGURE 3.** The taxonomy of algorithms for structured learning.

For the data representation in activity recognition and ADL classification scenarios, the bag-of-words (BoW) approach has proven to be convenient and successful. Originating in natural language processing, the BoW approach represents a text (such as a sentence or a document) as the bag (multiset) of its words, disregarding grammar and even word order but preserving multiplicity. An analogous bag-of-visual-words also has been successfully used for general image classification [79] and later for human action recognition and classification in video sequences [80]. Recent studies on human activity recognition show that the BoW representation allows achievement of high-performance action recognition [81], [82].

### Generative models

Generative models estimate the joint probability distribution of observation samples, which can be used to predict the most likely class to which a new sample belongs. They are called *generative*, because the model can be used to generate samples given the joint probability distribution. HMMs are a popular generative model that can deal with structured data where the IID assumption does not hold. In the context of traditional HMMs (having a finite number of discrete states), three important questions are asked as part of the model learning and its application on unseen data [83].

1) *Likelihood*: Given a model and a sequence of observations, how likely is it that this sequence was generated by the given model? The answer to this problem is given by the forward-backward algorithm.
2) *Decoding*: What is the most likely sequence of model states that generated a sequence of observations? The answer to this question is given by the Viterbi algorithm.
3) *Learning*: How should transition and emission probabilities be learned from observed sequences? The answer is given by the Baum–Welch algorithm, which can be seen as a special case of the expectation maximization algorithm and tries to optimize the model parameters to best describe the observation sequence, while using also the results of the two previous problems.

Because HMMs are suitable to model sequential data, it is a popular classification method in activity recognition. A variety of HMM-based variants is presented in a comprehensive survey by Turage et al. [84]. Also, the recognition of human motion data can be modeled with HMMs. Li [85] proposed a straightforward and effective motion descriptor based on oriented histograms of optical flow field sequences. Following dimensionality reduction performed by principal component analysis, the method was applied to human action recognition using the HMM approach. Yamato et al., in [86], used HMMs in their simplest form: training a set of HMMs, one for each action and modeling the observation probability function as a discrete distribution, adopting a mesh feature that computed frame by frame on the data [87]. Finally, the learning was based on the standard Baum–Welch approach. HMMs also can be applied on more complex data types, as demonstrated in Martinez et al. [88]. They proposed a framework for action recognition based on combining an HMM with a silhouette-based feature set. The proposed solution relies on a two-dimensional modeling of human actions based on motion templates, utilizing motion history images that combine viewpoint (spatial) and movement (temporal) representations.

Besides variations of the standard HMMs such as hierarchical HMMs (HHMMs) (where each class can be an HMM) and spectral HMMs (that can be used to do inference with unknown or continuous state spaces) the second important generative model are restricted Boltzmann machines (RBMs), which are implementing hidden layers in a stochastic neural network. They are often effective in cases where a lot of nonannotated data is available (typically thousands or tens of thousands of samples) but annotated data are scarce.

### Discriminative models

While the approaches based on HMMs discussed in the previous section have achieved unquestionable success in numerous applications, one of the drawbacks of HMMs is that they cannot take advantage of powerful discriminative learning techniques that have been developed for the classification of vectorial data, such as kernel machines or metric learning. Contrary to generative models, discriminative models do not try to model the underlying probability distribution, but instead estimate the conditional probability of the labeled sequences given the observations. Among other advantages, discriminative models are typically efficient at dealing with data in high dimensional spaces.

SVMs [89] are considered to be one of the most powerful discriminative classification methods that were applied to various problems including activity recognition and ADL classification [72]. By leveraging the structural risk minimization and the kernel trick concepts [90], SVMs are very effective in discriminating between classes even on very high dimensional vectorial data. SVMs are using only few samples (support vectors) to describe their decision boundaries making them memory efficient and resistant to noise or small class overlaps. The main drawback of SVMs is the high computational complexity of the training procedure that practically limits their applicability. Linear SVMs that are not using kernels are more attractive for large-scale data sets, however, they are not so powerful as SVMs that employ kernels.

Random forest is an ensemble learning method that is using a set of decision trees to solve classification tasks. Introduced by Breiman [91], random forest is famous for its ability to accurately separate data while being able to naturally handle both numerical and categorical features. Random forest requires no feature scaling procedure and by its intrinsic properties is not prone to overfitting. The employment of random forest in ADL classification scenarios can be found in [92].

CRFs have been proposed as a way to tackle similar structural problems where HMMs are applied, but relaxing certain of their assumptions [76], [93]. CRFs typically require less data to train than HMMs for a given performance level. They allow for the relaxation of the strong independence assumption between predictors and thereby allow for a richer set of features that can be partially overlapping. Their disadvantage is that they are computationally more complex (especially during training time) and, as all discriminative methods, they cannot make explicit estimations regarding the distribution of the observed variables and therefore cannot be used to sample from the learned model.

The standard CRF models the transitions between labels, thus capturing extrinsic dynamics, but lack the ability to represent internal substructure. Several modifications of the standard CRF have been proposed and applied to ADL classification including LDCRFs and semi-CRFs which, however, can be very computationally expensive.

**To address the drawbacks of generative and discriminative models, various studies have proposed to combine them into hybrid ones.**

### Hybrid models

To address the drawbacks of generative and discriminative models, various studies have proposed to combine them into hybrid ones. Hybrid models allow to leverage the ability to separate structural objects by learning similarity between them and, at the same time, to have access to all tools and advantages of generative models listed above. The Fisher kernel [94] and tangent vector of posterior log-odds (TOP) kernel are cases of such a hybrid method that relies on kernel metric distances. Other hybrid algorithms use a discriminant model such as an ANN or random forest to calculate the frame posterior probabilities and/or additional synthetic features while an HMM is responsible for modeling time dependencies on a metalevel through temporal smoothing on the estimated outcomes of the discriminative methods. Finally, CRF-HMM is another hybrid model often used in natural language processing due to its ability to model non-IID data both at the low as well as the high level of representation, for example to capture relations between letters and words.

Recently, hybrid HMM models have been also proposed for activity recognition. Ellis et al. [95] proposed to first learn low-level code-book representations for each sensor and use an random forest classifier to produce minute-level probabilities for each activity class. Subsequently, a higher-level HMM layer is used to learn patterns of transitions and durations of activities over time to smooth the minute-level predictions. Fisher kernel learning (FKL) is another approach that combines the flexibility of generative methods and the power of discriminative ones [96]. Fisher kernel representations have recently been applied to activity recognition problems [97]. The key intuition behind the Fisher kernel is that similar objects induce similar log-likelihood gradients in the parameters of a generative model allowing effective discrimination of these objects. To construct a Fisher kernel for structured objects, it is required to calculate the log-likelihood gradient for each of the objects in the parameters of a generative model. The Fisher kernel function can then be derived as a weighted inner product between the gradients of two structured objects [96]. The weighting is typically performed using the Fisher information metric; this weighting is necessary because different types of model parameters have different scales. Jaakkola and Haussler [94] have shown that the Fisher information metric is asymptotically immaterial, which is the theoretical basis for often assuming it to be an identity matrix. In this case, a (normalized) kernel is used that simply embeds objects in an Euclidean space by using the gradients induced by the objects as features [96].

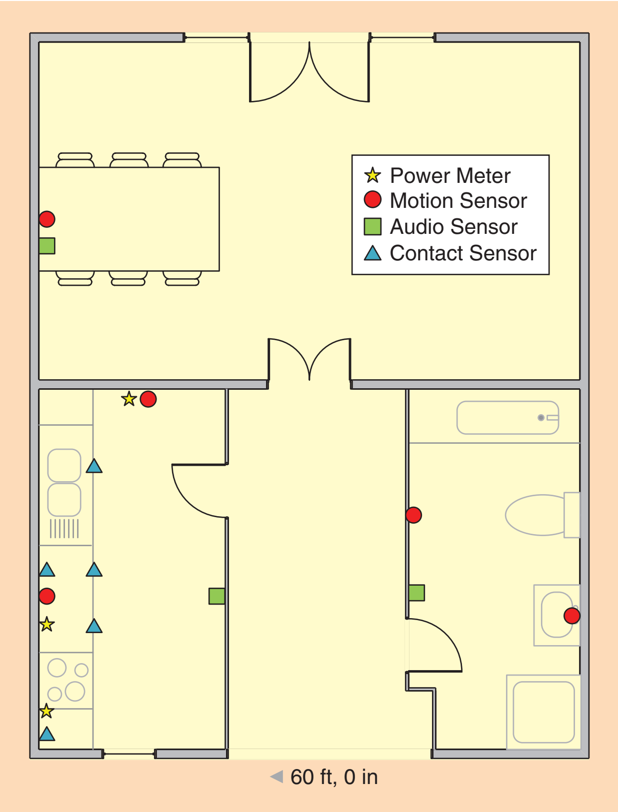


**FIGURE 4.** The layout and sensor setup of the first household.

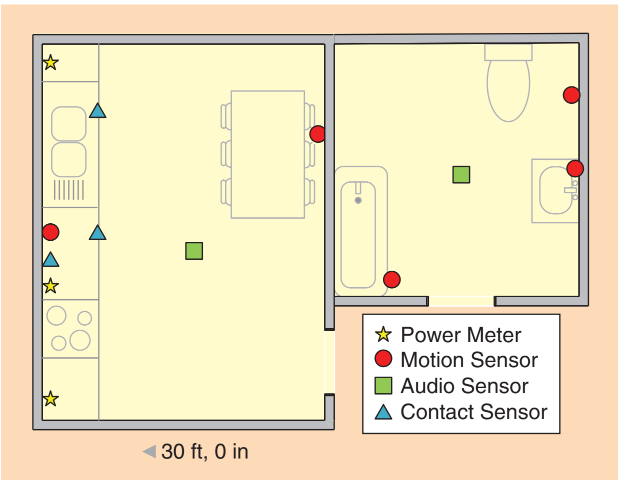


**FIGURE 5.** The layout and sensor setup of the second household.

## Experimental results

To evaluate different ADL classification methods, we consider the following three data sets:

- *Data set 1*: This data set [24] describes the activities of a 26-year-old man in his apartment where he lived alone. Fourteen state-change sensors were installed at doors, cupboards, the refrigerator, and the toilet flush. Sensors were left unattended, collecting data for 28 days in the apartment, resulting in 2,120 sensor events and 245 activity instances [24]. Seven ADL types were annotated.
- *Data set 2*: The second data set [98] is one of the multiple data sets recorded from the CASAS group of Washington State University. This particular data set was selected because it has the interesting property of capturing the activities of two people in a house, which can occur also in a practical application where one of the residents is a patient while the other is still healthy. The data were recorded over a period of two months using 34 sensors of four types: motion, item, door, and water sensors. Based on the annotations, 124 instances of activities of interest were captured.
- *Data set 3*: Finally, we recorded our own data set in two independent households, each for one week. The sensors installed include contact, motion, acoustic sensors, and power meters. In the first household, there were 7,856 events and in the second household there were 8,618 events, which resulted in 394 and 644 activity instances, respectively, that were annotated as five different ADL types. In the first household, five contact sensors, five motion sensors, three acoustic sensors, and three power meters were distributed over three rooms. In the second household, three contact sensors, five motion sensors, two acoustic sensors, and three power meters were distributed over two rooms. The room layout and sensor positions are shown in Figures 4 and 5.

The following four classifiers are most prominent in the community and are tested on all three data sets: 1) SVM, 2) random forest, 3) HMM, and 4) FKL. Regarding the evaluation of classification accuracy, the time slice accuracy is an established way of evaluating time series and represents the percentage of correctly classified frames, independently of the ADLs. However, since some of the ADLs have significantly larger duration compared to others, we also use the average class accuracy as a second criterion to avoid skewing the performance evaluation exclusively toward the dominant classes. To compute these accuracies, cross-validation is used. However, since the problem is structured learning and the IID assumption does not hold, the cross-validation is done at the event level, that is, without partitioning the samples of the same ADL. For example, of a total of

**Table 4. Time-slice and class-average accuracy for SVM, RF, HMM, and FKL on data set 1.**

| | | SVM | RF | HMM | FKL |
|---|---|---|---|---|---|
| Time-slice Accuracy | Idle | **0.82** | 0.80 | 0.19 | 0.67 |
| | Toilet | 0.58 | 0.81 | **0.87** | 0.77 |
| | Shower | 0.14 | 0.10 | 0.71 | **0.82** |
| | Bed | 0.37 | **0.99** | 0.96 | 0.95 |
| | Eat | 0.28 | 0.29 | 0.34 | **0.61** |
| | Drink | 0.38 | 0.13 | 0.00 | **0.42** |
| Class-average accuracy | | 0.43 | 0.52 | 0.51 | **0.71** |

30 food preparation instances, 20 might end up in the training and the remaining ten in the validation set, both with their full duration (all frames).

### *Data set 1*

The evaluation results on data set 1 are presented in Table 4. Of the different classifiers, FKL performs best on average. SVM has the lowest class-average accuracy since it does not consider time dependencies between data samples. The same holds for random forest, however it is very effective at modeling the variations in the execution of different ADLs and also is a bit less prone to overfit the dominant class (especially given that the annotation of ADLs is to some extent subjective resulting in some noise on the labels). The number of states per class for HMM was taken the same for all classes and this resulted in relatively low accuracy (in other words, this basic type of HMM cannot effectively model sequences that vary significantly in complexity for the different classes). The intrinsic number of states is different for various ADLs and should be determined carefully to obtain better performance. For example, there is a multitude of ways for doing food preparation but fewer ways for drinking. However, this would increase the complexity of HMM, requiring more training data and/or a model that can better capture hierarchy, like an HHMM. A full overview on the confusion between classes for FKL is provided in Figure 6.

Overall, the accuracies of classes Eat and, especially, Drink are substantially lower than those of the other classes because they are confused with each other and Idle. Both activities take place in the same location and involve similar actions such as opening a refrigerator or cupboard. The magnetic contact sensors that comprise the majority of the sensors in this data set are not very efficient for discriminating between events that are so similar that they involve triggering of the exact same sensors.

### *Data set 2*

Because the activities of two residents are recorded in this data set and the classifiers do not allow for overlapping activities, the classifiers are trained and tested on the activities of one of the residents. However, this means the learned activities in the training data as well as the activities in the test set can be

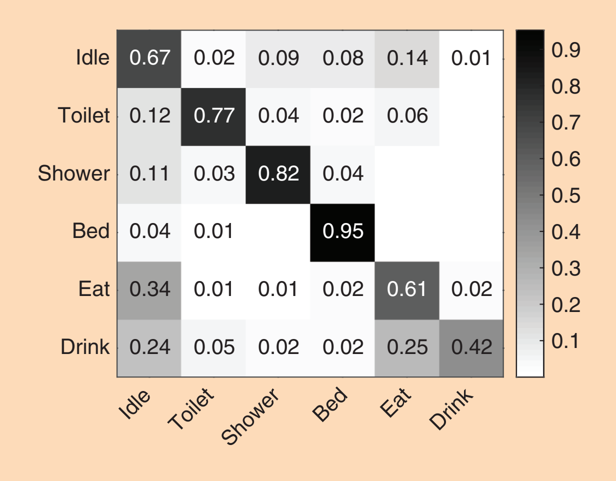


**FIGURE 6.** The confusion matrix for FKL on data set 1.

**Table 5. Time-slice and class-average accuracy for SVM, RF, HMM, and FKL on data set 2.**

| | | SVM | RF | HMM | FKL |
|---|---|---|---|---|---|
| Time-slice Acc. | Personal hygiene | 0.56 | 0.61 | 0.51 | **0.64** |
| | Sleep | 0.43 | 0.47 | 0.42 | **0.52** |
| | Work | 0.76 | 0.78 | **0.80** | 0.77 |
| | Meal preparation | 0.84 | 0.89 | **0.93** | 0.84 |
| | Watch TV | 0.30 | 0.35 | 0.00 | **0.41** |
| Class-average accuracy | | 0.58 | 0.62 | 0.53 | **0.64** |

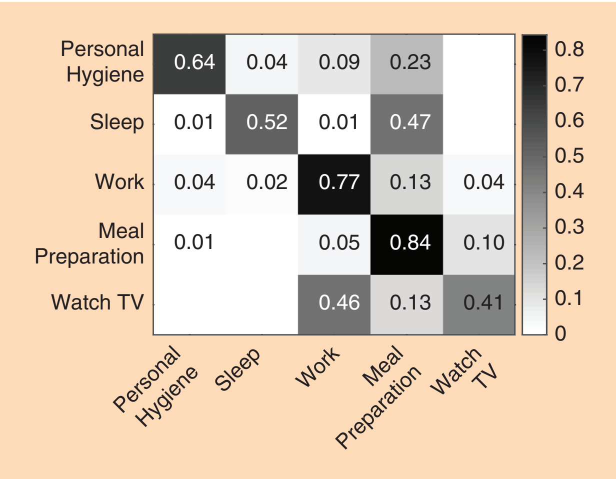


**FIGURE 7.** The confusion matrix for FKL on data set 2.

contaminated by the second resident. However, the evaluation results in Table 5 show that the classification algorithms can deal with this possible confusion to some extend, except for the Watch TV class. The effect of the contamination will reduce as the number of training examples increases, while the noise on the performance evaluation can be reduced by a bigger test set.

Similar to data set 1 with a one-person household, FKL shows the highest performance in terms of average accuracy. random forest and SVM perform slightly worse than FKL, while HMM provides the worst average accuracy, mostly because of the complete failure on the Watch TV class. A confusion matrix for FKL is provided in Figure 7. The classes Work and Meal Preparation can be discriminated the best among all other ADLs in this data set, while Personal Hygiene and Sleep demonstrate moderate recognition accuracy, and Watch TV performs the worst. Since two people are living in the apartment and can perform different activities at the same time and at different locations, dissimilar activities such as Sleep and Meal Preparation can be also confused.

**Automatic classification of ADLs enables automatic monitoring of the ability of an elderly person to live independently in his house and can allow for early detection of diseases such as Alzheimer and dementia.**

### *Data set 3*

For the third data set, we evaluated two scenarios. In the first scenario (one household) training and testing were done only using events coming from a single household. In the second scenario (two households) training and testing were performed on data taken from both households. All events from both households were shuffled in a way that the event sequence within one ADL is kept to preserve the structure in the data (non-IID assumption). Evaluation results on data set 3 for random forest, SVM, HMM, and FKL on one and two households are provided in Tables 6 and 7, respectively.

Also, for this data set, FKL outperforms the other methods, both for the one household and the two households scenario (see Tables 6 and 7). However, the performance of random forest and HMM are very close to FKL for the two households scenario. The HMM even has an improved performance compared to its performance for the one household scenario, probably because (for this data set) the improvement that it can gain with more data is bigger than the potential loss because of increased variation in the class representation.

The confusion matrices of FKL for the two scenarios are shown in Figures 8 and 9. In the one household scenario, FKL can achieve high performance, with the only noteworthy confusion appearing for very similar activities (see Figure 8). For example, Food Preparation and No ADL (which is defined as an activity that is not described by a formal ADL, e.g., standing in the kitchen, reading at a table, etc.) have an overlap while Continence and Hygiene also are partly difficult to discriminate in the absence of clear signature detections from the sensors on

**Table 6. Time-slice and class-average accuracy for SVM, RF, HMM, and FKL on data set 3 for the one household scenario.**

| | | SVM | RF | HMM | FKL |
|---|---|---|---|---|---|
| Time-slice Acc. | No ADL | 0.12 | 0.31 | 0.17 | **0.65** |
| | Continence | 0.39 | **0.90** | 0.75 | 0.70 |
| | Hygiene | **0.93** | 0.81 | 0.63 | 0.79 |
| | Showering | 0.89 | 0.73 | 0.55 | **0.94** |
| | Food preparation | 0.96 | 0.93 | **0.98** | 0.75 |
| Class-average accuracy | | 0.66 | 0.74 | 0.62 | **0.77** |

**Table 7. Time-slice and class-average accuracy for SVM, RF, HMM, and FKL on data set 3 for the two households scenario.**

| | | SVM | RF | HMM | FKL |
|---|---|---|---|---|---|
| Time-slice Acc. | No ADL | 0.31 | 0.19 | 0.62 | 0.63 |
| | Continence | 0.87 | 0.91 | 0.64 | 0.54 |
| | Hygiene | 0.72 | 0.74 | 0.57 | 0.78 |
| | Showering | 0.64 | 0.44 | 0.80 | 0.87 |
| | Food preparation | 0.88 | 0.94 | 0.84 | 0.74 |
| Class-average accuracy | | 0.68 | 0.64 | 0.69 | 0.71 |

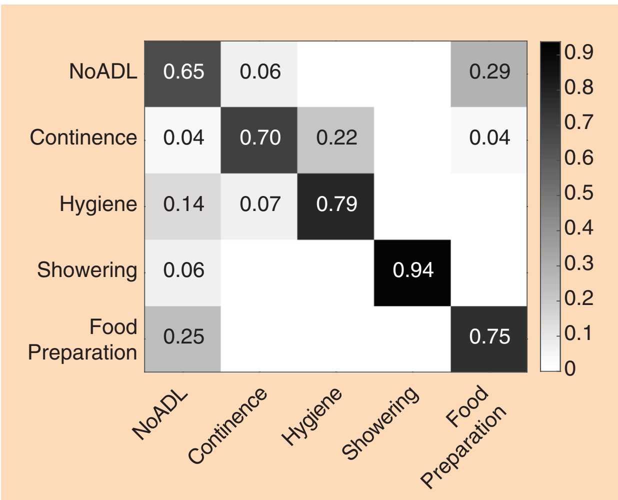


**FIGURE 8.** The confusion matrix for FKL on data set 3 for one household.

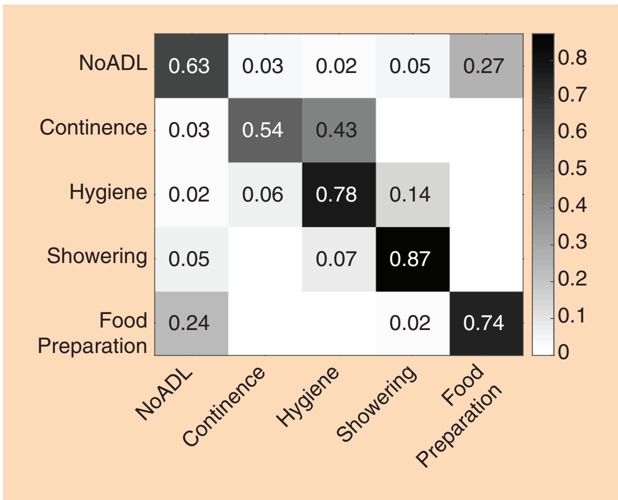


**FIGURE 9.** The confusion matrix for FKL on data set 3 for two households.

which the events are based (such as water running in the sink, toilet flushing, etc.).

The experiment with two households is more challenging and also more interesting, first of all because the sensor setups between the households are similar but not identical, and a BoW approach is used to map them in a common feature space for learning and inference. Furthermore, the households have different layouts and the persons are doing certain activities in very different ways. For example, in one of the households, cooking always involves opening and closing multiple kitchen cupboards while in the other household this is only done sporadically. Nevertheless, the classification scheme based on Fisher kernel is able to do learning and inference in the joined space, with the only significant overlap between classes appearing between the Continence and Hygiene, which is to some extent attributable to the different sensor setups and bath layouts.

**There is still progress to be made such that research in ADL classification can be reliably applied in a practical solution.**

## Conclusions

Automatic classification of ADLs is a crucial part of assisted living technologies. It enables automatic monitoring of the ability of an elderly person to live independently in his or her house and can allow for early detection of diseases such as Alzheimer's and dementia. ADL classification involves the whole chain from a plethora of wearable and nonwearable sensors, deployment options, and signal processing and machine-learning algorithms. Our study concludes that the recent developments in hybrid generative/discriminative methods, relying on kernel metric distances, are superior over traditional generative methods such as HMM and its variants. Specifically, FKL showed the best performance in a variety of data sets covering different activity types, sensors, and setups.

We expect continuing improvements on all aspects of the aforementioned chain ranging from improvement of existing sensor technologies, addition of new sensors, the acceptability of certain technologies up to various algorithmic aspects such as the generalizability and adaptiveness that are briefly detailed next.

Sensor technologies require improvements in several directions including size, accuracy, energy efficiency, and reliability. Establishing a common communication protocol would allow to create an unified framework, providing a significant speed up in infrastructure deployment. Reusability of sensors from smart home applications would be another boosting factor helping to reduce infrastructure and installation costs. Several legal issues related to the data ownership and data security have to be addressed to get acceptance for using ADL monitor systems.

While sensor technology improves, leading to higher-quality measurements and lower costs and maintenance, in practical applications the (elderly) user needs to be taken into account as well (see Figure 1). Specifically, user-centered design and transparency can help to increase the acceptance of users of technology in their home and their perceived privacy [15]. Furthermore, users are more and more exposed to sensor technology in other aspects of their lives, increasing their understanding and, thereby, acceptance.

Finally, we note that there is still progress to be made such that research in ADL classification can be reliably applied in a practical solution. This covers the optimization of sensor setups for cost effectiveness, adaptive classification algorithms that allow tracking changing behavior over time and robustness with respect to context changes such as handling of visitors, caregiving personnel, or pets. One of the biggest challenges from the signal processing and machine-learning side remains the generalizability over households. While training for an individual household is easily possible in a lab setup, this approach is not scalable to a real-world scenario with thousands of households and more. A successful approach for generalizability has to consider environmental/climate parameters, building layout, sensor placement, and the behavior of the elderly.

## Authors

***Christian Debes*** (cdebes@agtinternational.com) received B.Sc., M.Sc., and Ph.D. (with highest honors) degrees from from Technische Universität Darmstadt, Germany, in 2004, 2006, and 2010, respectively. He currently holds a position as lead data scientist at AGT International and lecturer at Technische Universität Darmstadt, Germany. He is a recipient of the IEEE Geoscience and Remote Sensing Society 2013 Data Fusion Best Classification Award and has authored more than 30 journal and conference papers in target detection, classification, and image processing. He is a member of the editorial board of *Digital Signal Processing* (Elsevier) and area editor of *IEEE Signal Processing Magazine*. He is a Senior Member of the IEEE.

***Andreas Merentitis*** (andreas.merentitis@ieee.org) received B.Sc., M.Sc., and Ph.D. degrees from National Kapodistrian University of Athens in 2003, 2005, and 2010 respectively. In 2011, he joined the research center of AGT International, Darmstadt, Germany, where he is now working as a senior data scientist. He has more than 30 publications related to different aspects of data analysis and machine learning. He was awarded an IEEE Certificate of Appreciation as a core member of the team that won the first place in the Best Classification Challenge of the 2013 IEEE Geoscience and Remote Sensing Society Data Fusion Contest. He is an associate editor of *IEEE Signal Processing Magazine* and a Member of the IEEE.

***Sergey Sukhanov*** (ssukhanov@agtinternational.com) received B.Sc. and M.Sc. degrees from Saint-Petersburg State Electrotechnical University, Russia, and Technische Universität Darmstadt, Germany, respectively. Since 2015, he has been working as a data scientist at AGT International, Darmstadt, Germany. His research interests include data analysis and machine learning. He is a Student Member of the IEEE.

***Maria Niessen*** (mniessen@agtinternational.com) received M.Sc. and Ph.D. degrees from the University of Groningen, The Netherlands. After two years of postdoctoral research at INCAS, she joined AGT International, Darmstadt, Germany, in 2012, where she is currently a senior data scientist. Her research interests include data analysis and acoustics.

***Nikolaos Frangiadakis*** (nfrangiadakis@agtinternational.com) received M.Sc. and Ph.D. degrees from the University of Maryland, College Park. Since 2011, he has been working as a senior researcher at AGT International, Darmstadt, Germany. He is recipient of the IEEE Geoscience and Remote Sensing Society 2013 Data Fusion Best Classification Award, and his research interests include data analysis, mobility modeling, simulation, and prediction. He is a Member of the IEEE.

***Alexander Bauer*** (abauer@agtinternational.com) received the Dipl.-Ing. degree from Karlsruhe Institute of Technology, Germany. He received his Ph.D. degree in computer science from Karlsruhe Institute of Technology, Germany, for his work on recommendation algorithms for satellite image interpretation systems of the German Strategic Intelligence Command. He is a lead data scientist at AGT International, Darmstadt, Germany, with ten years of experience in pattern recognition and machine learning. Since joining AGT in 2011, he designs data analytics components for anomaly detection, behavior analytics, and pattern learning.

## References

[1] M. M. Muszynska and R. Rau, "The old-age healthy dependency ratio in Europe," *J. Popul. Ageing*, vol. 5, no. 3, pp. 151–162, 2012.

[2] K. Giannakouris, "Ageing characterises the demographic perspectives of the Europe societies," in *Eurostat: Statistics in Focus*, vol. 72. 2008. [Online]. Available: http://bookshop.europa.eu/en/ageing-characterises-the-demographic-perspectives-of-the-european-societies-pbKSSF08072/

[3] G. K. Vincent and V. A. Velkoff, *The Next Four Decades: The Older Population in the United States: 2010 to 2050*, US Department of Commerce, Economics and Statistics Administration, U.S. Census Bureau, 2010.

[4] N. Farber and D. Shinkle, "Aging in place: A state survey of livability policies and practices," AARP Research Rep., Dec. 2011. [Online]. Available: https://assets.aarp.org/rgcenter/ppi/liv-com/aging-in-place-2011-full.pdf

[5] E. Oppers, K. Chikada, F. Eich, P. Imam, J. Kiff, M. Kisser, M. Soto, and T. Sun, "The financial impact of longevity risk," in *Global Financial Stability Report*. Washington, D.C.: International Monetary Fund, 2012.

[6] D. Silcock and D. Sinclair, "The cost of our ageing society," Tech. Rep., The International Longevity Centre, 2012. [Online]. Available: http://socialwelfare.bl.uk/subject-areas/services-client-groups/older-adults/ilcuk/141789The_cost_of_our_ageing_society.pdf

[7] D. J. Cook, J. C. Augusto, and V. R. Jakkula, "Ambient intelligence: Technologies, applications, and opportunities," *Pervasive Mob Comput*., vol. 5, no. 4, pp. 277–298, 2009.

[8] P. Rashidi and A. Mihailidis, "A survey on ambient-assisted living tools for older adults," *IEEE J. Biomed. Health Inform*., vol. 17, no. 3, pp. 579–590, 2013.

[9] G. Acampora, D. J. Cook, P. Rashidi, and A. V. Vasilakos, "A survey on ambient intelligence in healthcare," *Proc. IEEE*, vol. 101, no. 12, pp. 2470–2494, 2013.

[10] S. Katz, T. D. Downs, H. R. Cash, and R. C. Grotz, "Progress in development of the index of ADL," *Gerontologist*, vol. 10, no. 1, pp. 20–30, Spring 1970. [Online]. Available: http://gerontologist.oxfordjournals.org/content/10/1_Part_1/20.abstract

[11] L. Nygard, "Instrumental activities of daily living: A stepping-stone towards Alzheimer's disease diagnosis in subjects with mild cognitive impairment?," *Acta Neurol. Scand*., vol. 107, pp. 42–46, Feb. 2003.

[12] D. Galasko, D. Bennett, M. Sano, C. Ernesto, R. Thomas, M. Grundman, and S. Ferris, "An inventory to assess activities of daily living for clinical trials in Alzheimer's disease," *Alzheimer Dis. Assoc. Disord*., vol. 11, no. 2, pp. 33–39, 1997.

[13] R. S. Bucks, D. L. Ashworth, G. K. Wilcock, and K. Siegfried, "Assessment of activities of daily living in dementia: Development of the Bristol activities of daily living scale," *Age Ageing*, vol. 25, no. 2, pp. 113–120, 1996.

[14] P. Barberger-Gateau, D. Commenges, M. Gagnon, L. Letenneur, C. Sauvel, and J.F. Dartigues, "Instrumental activities of daily living as a screening tool for cognitive impairment and dementia in elderly community dwellers," *J. Amer. Geriatr. Soc*., vol. 40, no. 11, pp. 1129–1134, 1992.

[15] L. Lorenzen-Huber, M. Boutain, L. J. Camp, K. Shankar, and K. H. Connelly, "Privacy, technology, and aging: A proposed framework," *Ageing Int*., vol. 36, no. 2, pp. 232–252, 2011.

[16] W. L. Zagler, P. Panek, and M. Rauhala, *Ambient Assisted Living Systems-The Conflicts between Technology, Acceptance, Ethics and Privacy*. Wadern, Germany: Internat. Begegnungs-und Forschungszentrum für Informatik, 2008.

[17] P. Rashidi, D. J. Cook, L. B. Holder, and M. Schmitter-Edgecombe, "Discovering activities to recognize and track in a smart environment," *IEEE Trans. Knowl. Data Eng*., vol. 23, no. 4, pp. 527–539, 2011.

[18] S. Patel, H. Park, P. Bonato, L. Chan, and M. Rodgers, "A review of wearable sensors and systems with application in rehabilitation," *J. Neuroeng. Rehabil*., vol. 9, no. 12, pp. 1–17, 2012.

[19] M. Chan, E. Campo, and D. Esteve, "Assessment of activity of elderly people using a home monitoring system," *Int. J. Rehabil. Res*., vol. 28, no. 1, pp. 69–76, 2005.

[20] R. Suzuki, S. Otake, T. Izutsu, M. Yoshida, and T. Iwaya, "Rhythm of daily living and detection of atypical days for elderly people living alone as determined with a monitoring system," *J. Telemed. Telecare*, vol. 12, no. 4, pp. 208–214, 2006.

[21] M. Alwan, J. Leachtenauer, S. Dalal, D. Mack, S. Kell, and B. Turner, "Impact of monitoring technology in assisted living: Outcome pilot," *IEEE Trans. Inform. Technol. Biomed*., vol. 10, no. 12, pp. 192–198, 2006.

[22] H. W. Tyrer, M. A. Aud, G. Alexander, M. Skubic, and M. Rantz, "Early detection of health changes in older adults," in *Proc. Engineering in Medicine and Biology Society*, 2007, pp. 4045–4048.

[23] G. Virone, M. Alwan, S. Dalal, S. W. Kell, B. Turner, J. A. Stankovic, and R. Felder, "Behavioral patterns of older adults in assisted living," *IEEE Trans. Inform. Technol. Biomed*., vol. 12, no. 3, pp. 387–398, 2008.

[24] T. L. M. Van Kasteren, A. Noulas, G. Englebienne, and B. J. A. Kröse, "Accurate activity recognition in a home setting," in *Proc. ACM Int. Joint Conf. Pervasive and Ubiquitous Computing*, 2008, pp. 1–9.

[25] M. Skubic, G. Alexander, M. Popescu, M. Rantz, and J. Keller, "A smart home application to elder care: Current status and lessons learned," *Technol. Health Care*, vol. 17, no. 3, pp. 183–201, 2009.

[26] J. Biswas, S. D. Kumar, Q. Qui, V. C. Saradhi, and V. T. Pham, "Quality aware elderly people monitoring using ultrasonic sensors," in *Proc. 3rd Int. Conf. Smart Homes and Health Telematics*, 2005, pp. 107–115.

[27] A. Fleury, M. Vacher, and N. Noury, "SVM-based multimodal classification of activities of daily living in health smart homes: Sensors, algorithms, and first experimental results," *IEEE Trans. Inform. Technol. Biomed*., vol. 14, no. 2, pp. 274–283, 2010.

[28] L. Vuegen, B. Van Den Broeck, P. Karsmakers, H. Van Hamme, and B. Vanrumste, "Automatic monitoring of activities of daily living based on real-life acoustic sensor data: A preliminary study," in *Proc. Workshop on Speech and Language Processing for Assistive Technologies*, 2013, vol. 4, pp. 113–118.

[29] J. Biswas, F. Naumann, and Q. Qiu, "Assessing the completeness of sensor data," in *Proc. Database Systems for Advanced Applications*, 2006, pp. 717–732.

[30] V. T. Pham, Q. Qiu, A. A. P. Wai, and J. Biswas, "Application of ultrasonic sensors in a smart environment," *Pervasive Mob. Comput.*, vol. 3, no. 2, pp. 180–207, 2007.

[31] K. Doughty and J. Costa, "Continuous automated telecare assessment of the elderly," *J. Telemed. Telecare*, vol. 3, no. 1, pp. 23–25, 1997.

[32] K. Z. Haigh, L. M. Kiff, and G. Ho, "The independent lifestyle assistant: Lessons learned," *Assist. Technol.*, vol. 18, no. 1, pp. 87–106, 2006.

[33] K. Cameron, K. Hughes, and K. Doughty, "Reducing fall incidence in community elders by telecare using predictive systems," in *Proc. 19th Annu. Int. Conf. IEEE Engineering in Medicine and Biology Society*, 1997, vol. 3, pp. 1036–1039.

[34] E. M. Tapia, S. S. Intille, and K. Larson, "Activity recognition in the home using simple and ubiquitous sensors," *Pervasive Comput.*, vol. 3001, pp. 158–175, Apr. 2004.

[35] J. Han and B. Bhanu, "Human activity recognition in thermal infrared imagery," in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, 2005, pp. 17–17.

[36] P. Hevesi, S. Wille, G. Pirkl, N. Wehn, and P. Lukowicz, "Monitoring household activities and user location with a cheap, unobtrusive thermal sensor array," in *Proc. ACM Int. Joint Conf. Pervasive and Ubiquitous Computing*, 2014, pp. 141–145.

[37] Y. Zigel, D. Litvak, and I. Gannot, "A method for automatic fall detection of elderly people using floor vibrations and sound - proof of concept on human mimicking doll falls," *IEEE Trans. Biomed. Eng.*, vol. 56, no. 12, pp. 2858–2867, 2009.

[38] S. S. Intille, K. Larson, J. S. Beaudin, J. Nawyn, E. M. Tapia, and P. Kaushik, "A living laboratory for the design and evaluation of ubiquitous computing technologies," in *Proc. CHI'05 Extended Abstracts on Human Factors in Computing Systems*, 2005, pp. 1941–1944.

[39] T. Tsukiyama, "Ambient sensor system for in-home health monitoring," in *Proc. The 4th Int. Conf. Ambient Computing, Applications, Services and Technologies*, 2014, pp. 47–50.

[40] L. Hu, Y. Chen, S. Wang, and L. Jia, "A nonintrusive and single-point infrastructure-mediated sensing approach for water-use activity recognition," in *Proc. IEEE Int. Conf. Embedded and Ubiquitous Computing*, 2013, pp. 2120–2126.

[41] J. H. Lim, H. Jang, J. Jang, and P. Soo-Jun, "Daily activity recognition system for the elderly using pressure sensors," in *Proc. Engineering in Medicine and Biology Society*, 2008, pp. 5188–5191.

[42] N. Noury, T. Herve, V. Rialle, G. Virone, E. Mercier, G. Morey, A. Moro, and T. Porcheron, "Monitoring behavior in home using a smart fall sensor and position sensors," in *Proc. Int. Conf. Microtechnologies in Medicine and Biology*, 2000, pp. 607–610.

[43] A. Glascock and D. Kutzik, "Behavioral telemedicine: A new approach to the continuous nonintrusive monitoring of activities of daily living," *Telemed. J.*, vol. 6, no. 1, pp. 33–34, 2004.

[44] L. Atallah, B. Lo, R. Ali, R. King, and Y. Guang-Zhong, "Real-time activity classification using ambient and wearable sensors," *IEEE Trans. Inform. Technol. Biomed.*, vol. 13, no. 6, pp. 1031–1039, 2009.

[45] A. Fleury, N. Noury, M. Vacher, H. Glasson, and J. F. Seri, "Sound and speech detection and classification in a health smart home," in *Proc. 30th Annu. Int. Conf. Engineering in Medicine and Biology Society*, 2008, pp. 4644–4647.

[46] M. Popescu, Y. Li, M. Skubic, and M. Rantz, "An acoustic fall detector system that uses sound height information to reduce the false alarm rate," in *Proc. 30th Annu. Int. Conf. Engineering in Medicine and Biology Society*, 2008, pp. 4628–4631.

[47] T. D. Hunt, D. Rajendran, M. Nikora, S. Bennett, and A. Fendall, "A minimally intrusive monitoring system that utilizes electricity consumption as a proxy for wellbeing," *J. Appl. Comput. Inform. Technol.*, vol. 18, no. 2, 2014. [Online]. Available: https://doaj.org/article/30741addf579403387b4bb9d3fb2bff3

[48] G. C. Franco, F. Gallay, M. Berenguer, C. Mourrain, and P. Couturier, "Noninvasive monitoring of the activities of daily living of elderly people at home—A pilot study of the usage of domestic appliances," *J. Telemed. Telecare*, vol. 14, no. 5, pp. 231–235, 2008.

[49] D. Fei and J.-Y. Xiong, "The investigation of the elder's monitoring system based on life supplying line," in *Proc. IEEE Int. Conf. Industrial Technology*, 2005, pp. 314–318.

[50] T. Huỳnh, U. Blanke, and B. Schiele, "Scalable recognition of daily activities with wearable sensors," in *Proc. Location-and Context-Awareness*, 2007, pp. 50–67.

[51] Y. Yang, J. S. Wang, and J. P. Chen, "Using acceleration measurements for activity recognition: An effective learning algorithm for constructing neural classifiers," *Pattern Recognit. Lett.*, vol. 29, no. 16, pp. 2213–2220, 2008.

[52] A. M. Khan, Y. K. Lee, and S. Y. Lee, "Accelerometer's position free human activity recognition using a hierarchical recognition model," in *Proc. 12th IEEE Int. Conf. e-Health Networking Applications and Services*, 2010, pp. 296–301.

[53] H. Gjoreski, M. Lustrek, and M. Gams, "Accelerometer placement for posture recognition and fall detection," in *Proc. 7th Int. Conf. Intelligent Environments*, 2011, pp. 47–54.

[54] M. J. Mathie, A. Coster, N. Lovell, G. Celler, S. Lord, and A. Tiedemann, "A pilot study of long-term monitoring of human movements in the home using accelerometry," *J. Telemed. Telecare*, vol. 10, no. 3, pp. 144–151, 2004.

[55] Q. Li, J. A. Stankovic, M. A. Hanson, A. T. Barth, J. Lach, and G. Zhou, "Accurate, fast fall detection using gyroscopes and accelerometer-derived posture information," in *Proc. 6th Int. Workshop on Wearable and Implantable Body Sensor Networks*, 2009, pp. 138–143.

[56] R. Amstutz, O. Amft, B. French, A. Smailagic, D. Siewiorek, and G. Troster, "Performance analysis of an hmm-based gesture recognition using a wristwatch device," in *Proc. Int. Conf. Computational Science and Engineering*, 2009, pp. 303–309.

[57] J. Merilahti, J. Pärkkä, K. Antila, P. Paavilainen, E. Mattila, E. J. Malm, A. Saarinen, and I. Korhonen, "Compliance and technical feasibility of long-term health monitoring with wearable and ambient technologies," *J. Telemed. Telecare*, vol. 15, no. 6, pp. 302–309, 2009.

[58] K. Van Laerhoven, D. Kilian, and B. Schiele, "Using rhythm awareness in long-term activity recognition," in *Proc. IEEE Int. Symp. Wearable Computers*, 2008, pp. 63–66.

[59] T. Maekawa, Y. Kishino, Y. Sakurai, and T. Suyama, "Activity recognition with hand-worn magnetic sensors," *Pers. Ubiquitous Comput.*, vol. 17, no. 6, pp. 1085–1094, 2013.

[60] J. Lester, T. Choudhury, and G. Borriello, "A practical approach to recognizing physical activities," in *Proc. Pervasive Computing*, 2006, pp. 1–16.

[61] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, "Activity recognition using cell phone accelerometers," *SIGKDD Explor. Newsl.*, vol. 12, no. 2, pp. 74–82, 2010.

[62] S. Park and H. Kautz, "Hierarchical recognition of activities of daily living using multi-scale, multi-perspective vision and RFID," in *Proc. IET 4th Int. Conf. Intelligent Environments*, IET, 2008, pp. 1–4.

[63] M. Stikic, T. Huynh, K. Van Laerhoven, and B. Schiele, "ADL recognition based on the combination of RFID and accelerometer sensing," in *Proc. 2nd Int. Conf. Pervasive Computing Technologies for Healthcare*, 2008, pp. 258–263.

[64] M. Buettner, R. Prasad, M. Philipose, and D. Wetherall, "Recognizing daily activities with RFID-based sensors," in *Proc. 11th Int. Conf. Ubiquitous Computing*, ACM, 2009, pp. 51–60.

[65] M. Philipose, K. P. Fishkin, M. Perkowitz, D. J. Patterson, D. Fox, H. Kautz, and D. Hahnel, "Inferring activities from interactions with objects," *IEEE Pervasive Comput.*, vol. 3, no. 4, pp. 50–57, 2004.

[66] A. Matic, P. Mehta, J. M. Rehg, V. Osmani, and O. Mayora, "Monitoring dressing activity failures through RFID and video," *Methods Inform. Med.*, vol. 51, no. 1, pp. 45, 2012.

[67] S. Junnila, H. Kailanto, J. Merilahti, A.-M. Vainio, A. Vehkaoja, M. Zakrzewski, and J. Hyttinen, "Wireless, multipurpose in-home health monitoring platform: Two case trials," *IEEE Trans. Inform. Technol. Biomed.*, vol. 14, no. 2, pp. 447–455, 2010.

[68] A. Bulling, J. A. Ward, H. Gellersen, and G. Troster, "Eye movement analysis for activity recognition using electrooculography," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 4, pp. 741–753, 2011.

[69] M. Chan, D. Estéve, C. Escriba, and E. Campo, "A review of smart homes—present state and future challenges," *Comput. Methods Programs Biomed.*, vol. 91, no. 1, pp. 55–81, 2008.

[70] P. Rashidi and D. J. Cook, "Keeping the resident in the loop: Adapting the smart home to the user," *IEEE Trans. Syst. Man Cybern. Part A Syst. Hum.*, vol. 39, no. 5, pp. 949–959, 2009.

[71] M. J. Rantz, M. Skubic, S. J. Miller, C. Galambos, G. Alexander, J. Keller, and M. Popescu, "Sensor technology to support aging in place," *J. Amer. Med. Dir. Assoc.*, vol. 14, no. 6, pp. 386–391, 2013.

[72] S. Chernbumroong, S. Cang, A. Atkins, and H. Yu, "Elderly activities recognition and classification for applications in assisted living," *Expert Syst. Appl.*, vol. 40, no. 5, pp. 1662–1674, 2013.

[73] L. Gao, A. K. Bourke, and J. Nelson, "Evaluation of accelerometer-based multi-sensor versus single-sensor activity recognition systems," *Med. Eng. Phys.*, vol. 36, no. 6, pp. 779–785, June 2014.

[74] D. J. Patterson, D. Fox, H. Kautz, and M. Philipose, "Fine-grained activity recognition by aggregating abstract object usage," in *Proc. 9th IEEE Int. Symp. Wearable Computers*, 2005, pp. 44–51.

[75] R. Helaoui, M. Niepert, and H. Stuckenschmidt, "Recognizing interleaved and concurrent activities: A statistical-relational approach," in *Proc. IEEE Int. Conf. Pervasive Computing and Communications (PerCom)*, 2011, pp. 1–9.

[76] E. Nazerfard, B. Das, L. B. Holder, and D. J. Cook, "Conditional random fields for activity recognition in smart environments," in *Proc. ACM Int. Health Informatics Symp.*, 2010, pp. 282–286.

[77] G. Virone, N. Noury, and J. Demongeot, "A system for automatic measurement of circadian activity deviations in telemedicine," *IEEE Trans. Biomed. Eng.*, vol. 49, no. 12, pp. 1463–1469, Dec. 2002.

[78] C. Zhu and W. Sheng, "Motion- and location-based online human daily activity recognition," *Pervasive Mob. Comput.*, vol. 7, no. 2, pp. 256–269, 2011.

[79] J. Zhang, M. Marszalek, S. Lazebnik, and C. Schmid, "Local features and kernels for classification of texture and object categories: A comprehensive study," *Int. J. Comput. Vis.*, vol. 37, no. 2, pp. 213–238, 2007.

[80] H. Wang, M. M. Ullah, A. Klser, I. Laptev, and C. Schmid, "Evaluation of local spatiotemporal features for action recognition," in *Proc. British Machine Vision Conf.*, 2009, pp. 124.1–124.11.

[81] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld, "Learning realistic human actions from movies," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2008, pp. 1–8.

[82] P. Bilinski and F. Bremond, "Evaluation of local descriptors for action recognition in videos," *Comput. Vis. Syst. Springer*, vol. 6962, pp. 61–70, 2011. [Online]. Available: http://link.springer.com/chapter/10.1007%2F978-3-642-23968-7_7#page-1

[83] L. Rabiner and B. H. Juang, "An introduction to hidden Markov models," *IEEE ASSP Mag.*, vol. 3, no. 1, pp. 4–16, Jan. 1986.

[84] P. Turaga, R. Chellappa, V. S. Subrahmanian, and O. Udrea, "Machine recognition of human activities: A survey," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 18, no. 11, pp. 1473–1488, Nov. 2008.

[85] X. Li, "HMM-based action recognition using oriented histograms of optical flow field," *Electron. Lett.*, vol. 43, no. 10, pp. 560–561, 2007.

[86] J. Yamato, J. Ohya, and K. Ishii, "Recognizing human action in time-sequential images using hidden Markov model," in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, June 1992, pp. 379–385.

[87] C. Yip, *Machine Recognition of Multi-font Printed Chinese Characters*. Pokfulam, Hong Kong: Univ. Hong Kong, 1990.

[88] F. Martinez-Contreras, C. Orrite-Urunuela, E. Herrero-Jaraba, H. Ragheb, and S. A. Velastin, "Recognizing human actions using silhouette-based HMM," in *Proc. 6th IEEE Int. Conf. Advanced Video and Signal Based Surveillance, 2009 (AVSS'09)*, Sept. 2009, pp. 43–48.

[89] V. N. Vapnik, *The Nature of Statistical Learning Theory*. New York: Springer-Verlag, 1995.

[90] B. Scholkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Cambridge, MA: MIT Press, 2001.

[91] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.

[92] P. Urwyler, L. Rampa, R. Stucki, M. Büchler, R. Müri, U. P Mosimann, and T. Nef, "Recognition of activities of daily living in healthy subjects using two ad-hoc classifiers," *Biomed. Eng. Online*, vol. 14, no. 1, 2015. [Online]. Available: http://link.springer.com/article/10.1186%2Fs12938-015-0050-4#/page-1

[93] K. Bousmalis, S. Zafeiriou, L.-P. Morency, and M. Pantic, "Infinite hidden conditional random fields for human behavior analysis," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 24, no. 1, pp. 170–177, Jan. 2013.

[94] T. S. Jaakkola and D. Haussler, "Exploiting generative models in discriminative classifiers," in *Proc. Conf. Advances in Neural Information Processing Systems II*, MIT Press, Cambridge, MA, USA, 1999, pp. 487–493.

[95] K. Ellis, J. Kerr, S. Godbole, and G. Lanckriet, "Multi-sensor physical activity recognition in free-living," in *Proc. 2014 ACM Int. Joint Conf. Pervasive and Ubiquitous Computing: Adjunct Publication, UbiComp'14 Adjunct*, New York, NY, USA, ACM, 2014, pp. 431–440.

[96] L. J. P. Van der Maaten, "Learning discriminative fisher kernels," in *Proc. 28th Int. Conf. Machine Learning*, 2011, pp. 217–224.

[97] I. Alexiou, T. Xiang, and S. Gong, "Learning a joint discriminative-generative model for action recognition," in *Proc. Int. Conf. Systems, Signals and Image Processing*, London, UK, 2015., pp. 1–4. [Online]. Available: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?reload=true&arnumber=7313922

[98] D. J. Cook and M. Schmitter-Edgecombe, "Assessing the quality of activities in a smart environment," *Methods Inform. Med.*, vol. 48, no. 5, pp. 480–485, 2009.

SP