Journal of Universal Computer Science, vol. 19, no. 9 (2013), 1295-1314
submitted: 12/2/13, accepted: 26/3/13, appeared: 1/5/13 © *J.UCS*

Energy Efficient Smartphone-Based Activity Recognition using Fixed-Point Arithmetic

Davide Anguita, Alessandro Ghio, Luca Oneto
(Università degli Studi di Genova, Genoa, Italy
{davide.anguita, alessandro.ghio, luca.oneto}@unige.it)

Xavier Parra, Jorge L. Reyes-Ortiz
(Universitat Politècnica de Catalunya, Vilanova i la Geltrú, Spain
xavier.parra@upc.edu, jorge.luis.reyes@estudiant.upc.edu)

Abstract: In this paper we propose a novel energy efficient approach for the recognition of human activities using smartphones as wearable sensing devices, targeting assisted living applications such as remote patient activity monitoring for the disabled and the elderly. The method exploits fixed-point arithmetic to propose a modified multiclass Support Vector Machine (SVM) learning algorithm, allowing to better preserve the smartphone battery lifetime with respect to the conventional floating-point based formulation while maintaining comparable system accuracy levels. Experiments show comparative results between this approach and the traditional SVM in terms of recognition performance and battery consumption, highlighting the advantages of the proposed method.

Key Words: Activity Recognition, Remote Monitoring, SVM, Smartphones, Energy Efficiency, Fixed-Point Arithmetic, Assisted Healthcare
Category: D.2, H.1.2, I.2, I.2.6, I.2.9, J.3, J.7

1 Introduction

Remote patient monitoring is nowadays allowing disabled and elderly patients a continuous health and well-being supervision while they perform regular activities throughout the day. Recent population benchmarks show that world population is aging rapidly. As an example, the projections of changes in population structure by main age groups in Europe are showing that by 2060 the elderly (namely people over 65 years) will be near 30% of its population [Eur, 2011]. This represents an alarming growth of more than 70% of this age group, bringing new challenges to the research community, which aims to find beneficial alternatives for ensuring healthy living to the people.

An extensive research has been particularly focused on Home Care Monitoring and the development of Smart Homes [Silva et al., 2012; José et al., 2010; García-Vázquez et al., 2010] for specific applications such as in elderly care and child care. In general, these novel households are arranged with multimodal technologies involving environmental sensors, user interfaces, computing devices and actuators which aim to guarantee supervision and fast response to people with frailty or chronic diseases such as Parkinson's disease (PD) and visual impairments.

Even though patients spend most of their time at home, they also commute from one place to another: for example, to buy groceries, to walk their dogs or to visit neighbors and friends. Home environmental sensors are limited by their infrastructure and cannot provide monitoring outside the house boundaries. Contrariwise, on-body sensors can improve and enlarge the range of operation of the patient monitoring task, not only by being capable of measuring a variety of body signals (e.g. physiological, motion, location) but also providing portable and off-site patient supervision. Unfortunately there are also limitations that arise with the use of body sensors such as patient discomfort while wearing them and energy-limited mobile devices. In this work we deal with these two issues by exploring the use of smartphones as monitoring devices for the classification of Activities of Daily Living (ADL).

Smartphones emerge from the integration of new services and features to mobile phones that complement the traditional telephony service (e.g. Internet access, gaming, location–based services and multisensing capabilities,etc.). They are playing an important role in the exploration of novel alternatives for the retrieval of information directly from the users. It is foreseen that these devices will be able to monitor and learn from our actions effectively and unobtrusively, and consequently assist us to better decide about our future behavior [Cook and Das, 2012]. Human Activity Recognition (HAR) is a research field that aims to identify the actions carried out by one or more subjects through the gathering and understanding of context information about the user state and its surrounding environment. This is done by the exploitation of environmental and on-body sensors, and distributed computing resources. Accelerometry is one of the mechanisms used for the retrieval of body motion information and which has been applied for the recognition of human activities [Allen et al., 2006]. One of the advantages of the recent smartphone technologies is that they are incorporating inertial sensors such as accelerometers, gyroscopes and magnetometers. These sensors were initially purposed for allowing enriched user interfaces and augmented gaming options, but they are now being exploited for HAR through the use of supervised Machine Learning (ML) approaches for a wide range of new applications that benefit from the phone's processing and opportunistic sensing capabilities.

In this paper, we employ smartphones for HAR targeting potential applications in areas such as healthcare and assisted living technologies. We concentrate efforts in dealing with energy efficiency, which is currently a limitation of these mass-marketed devices, and propose a novel approach that requires fewer system resources for its operation and aims to balance the trade-off between recognition accuracy and computational cost. For such purposes, we introduce the exploitation of the MultiClass Hardware-Friendly Support Vector Machine (MC-HF-SVM) approach, which makes use of fixed-point arithmetic for the recognition of activities instead of the conventionally used floating-point arithmetic algorithms. MC-HF-SVM allows to vary fixed-point number representation (number of bits) to control over model accuracy and complexity, leading to improvements in terms of both recognition accuracy and battery energy sparing. Being this contribution an extension of the research presented in [Anguita et al., 2012a], a demonstration of MC-HF-SVM benefits is also exposed on the subject of Statistical Learning Theory (SLT). In order to verify the effectiveness of the proposed approach, a collection of human physical activities (*standing, sitting, laying, walking, walking upstairs* and *walking downstairs*) were selected for classification and an experimental setup was arranged for data collection. Body motion signals were read from the smartphone embedded triaxial accelerometer while a group of volunteers performed the aforementioned activities, and then used for developing and testing the MC-HF-SVM approach, as well as for comparing MC-HF-SVM against conventional SVM-based techniques.

The paper has been structured in the following way: Section 2 presents a comprehensive collection of research literature including recent works on HAR, wearable systems and ML. Then in Section 3, the proposed HAR methodology is described from the experimental setup to the mathematical formulation of the MC-HF-SVM algorithm and its relationship with SLT. Moreover in Section 4, we show experimental results comparing the traditional SVM with the hardware-friendly approach in terms of system accuracy and energy consumption. Finally conclusions and future directions are portrayed in Section 5.

2 Related Work

Several approaches have been previously proposed in literature for the recognition of human activities covering diverse application domains such as healthcare, smart homes, ubiquitous computing, ambient assisted living, surveillance and security [Choudhury et al., 2008; Cedras and Shah, 1995; Turaga et al., 2008; Poppe, 2010]. These approaches can be categorized according to many different criteria: by *sensor type*, which is reliant on the signals measured (e.g. inertial, vision-based and physiological [Lara and Labrador, 2012b]); by *sensor location*, namely external sensing when sensors are located in fixed positions in the environment and wearable sensing when they are body-attached [Yang and Yacoub, 2006]; by *modeling principle*, which can be data- or knowledge-driven depending on whether the HAR models are built given pre-existing datasets or from the exploitation of prior knowledge regarding a particular domain [Chen et al., 2012a,b]; by *learning approach*, which can be either supervised, semi-supervised or unsupervised [Kwapisz et al., 2011; Stikic et al., 2011; Wyatt et al., 2005]. In this work we focus on a supervised smartphone-based HAR approach that make use of wearable inertial sensors (accelerometers) following a data-driven perspective.

Wearable systems have particularly grabbed the attention of the HAR research community [Bao and Intille, 2004; Lukowicz et al., 2004; Lee and Mase, 2002; Mantyjarvi et al., 2001] due to the ease of obtaining activity information (e.g. body motion, temperature and heart rate) directly from the user, unobtrusively and virtually at any location without the need of fixed infrastructure as opposed to vision-based systems. The work presented in [Bao and Intille, 2004] was pioneer in developing an approach for the classification of ADL using five body-worn accelerometers and employing well-known ML classifiers. Since then, other approaches have been also proposed [Lara and Labrador, 2012b] targeting different applications: for example, from the medical standpoint, monitoring systems have been presented for the detection of different attributes in elder

PD patients such as gait parameters, motion disorders and falls using on-body accelerometer [Sama et al., 2012; Herrlich et al., 2011].

More recently, research efforts have concentrated on exploiting smartphones for HAR. Smartphone-based applications offer various benefits when compared with other well-known wearable HAR alternatives that use special-purpose devices or body sensor networks attached to the body (e.g. [Mannini and Sabatini, 2010; Vinh et al., 2011]). Their main advantages relies on the easy device portability, the unobtrusive sensing provided by its embedded sensors and the processing power of nowadays devices that allow to perform online expensive computations all in one place. Several smartphone-based approaches that make use of various smartphone-embedded sensors (e.g. accelerometer, gyroscope, GPS, magnetometer) have been investigated [Mannini and Sabatini, 2010; Lara and Labrador, 2012b,a; Cook and Das, 2007; Berchtold et al., 2010; Kwapisz et al., 2011].

Machine Learning approaches that have been already applied for the recognition of activities include: Naive Bayes [Jatoba et al., 2008], and Markov chains [Mannini and Sabatini, 2010], Decision Trees [Maurer et al., 2006a], and Support Vector Machines (SVMs) [Vapnik, 1995]. Our approach exploits SVMs for the classification of activities similarly to other works which have successfully employed them [Maurer et al., 2006b; Khan et al., 2010; He and Jin, 2009; Ravi et al., 2005]. Furthermore, they have shown to be effective in heterogeneous types of recognition such as in handwritten characters [LeCun et al., 1995] and speech [Ganapathiraju et al., 2004]. SVMs have also been selected for this work because they provide a good compromise between accuracy and training time while also count with a variety of publicly available learning tools for experimentation such as LIBSVM [Chang and Lin, 2011]. However, while characterized by several appealing characteristics, one of the main drawbacks of SVM consists in its naïve two-class nature, that makes generalization to multiclass problems (as in the typical case of HAR) not straightforward.

Different approaches have been explored for targeting this issue [LeCun et al., 1995]. The two most commonly used methods are: One-Vs-All (OVA) and One-Vs-One (OVO), where each particular class is compared using a binary classifier against the rest of classes either all together (OVA) or one by one (OVO) to determine the most likely class for each new sample. In particular, we have selected the One-Vs-All (OVA) method and customized it to the fixed-point arithmetic case. The performance of the OVA approach is comparable to the OVO classification as it has previously been confirmed in [Rifkin and Klautau, 2004]. Moreover its produced model needs less memory when compared against the OVO method, bringing up an advantage taking into account the limited resources available.

Exploiting SVM models for HAR on smartphones, however, requires several floating-point operations to be carried out per second: despite not being an issue from a theoretical point of view, this could lead to battery discharge after few hours of continuous operation, making this approach unfeasible to allow people's mobility. In this work we will thus explore a fixed-point arithmetic based reformulation of the conventional SVM, targeted towards multiclass classification.

Up to date we have no knowledge of other research works that have incorporated fixed-point arithmetic into the learning algorithms for the classification of human activities. However, extensive research on fixed-point arithmetic has been developed to integrate ML models on hardware with limited resources (e.g. [Wawrzynek et al., 1993]). This was initially motivated because the assemble of devices with floating-point units was unavailable. Moreover, limited devices are usually preferred for specific-purpose applications if they demonstrate similar performance to traditional processing units as their production (and/or acquisition) costs are generally lower. Nowadays, it has become particularly interesting to retake these approaches and apply them in the development of software applications for portable devices such as smartphones which are highly demanding in terms of energy consumption and system resources management. The term Hardware-Friendly SVM (HF-SVM) was first presented in [Anguita et al., 2007]. This method was designed for binary classification problems by employing fixed-point arithmetic in the feed-forward phase of the SVM classifier, with the purpose of allowing its use in hardware-limited devices. In this work, we adapt the model for the multiclass problem targeted towards HAR on smartphones.

In the last decades several works have been devoted to adapt Machine Learning (ML) approaches to specific hardware platforms [Genov and Cauwenberghs, 2003; Lee et al., 2003; Irick et al., 2008; Epitropakis et al., 2010] and, in particular, to analyze the effects of parameter quantization on the training and feed-forward phases [Anguita et al., 2007; Lesser et al., 2011; Neven et al., 2009]. Motivations for these activities are usually linked to application-specific requirements but also to the basic principle of the Statistical Learning Theory (SLT) [Vapnik, 1995] where we have to search for the easiest model that correctly classifies the available data. The introduction of bit–based hypothesis spaces brings widespread benefits on the learning process of classifiers (i.e. classes of functions where models are described through a limited number of bits). This is due to the fact that reducing the number of bits largely influences the complexity of the hypothesis space [Anguita et al., 2013], which is a key issue in Machine Learning as underlined in [Shawe-Taylor et al., 1998; Bartlett et al., 2005]. If we are able to reduce the complexity of the hypothesis space without affecting the ability of the algorithm to learn the function with low empirical error, in practice, we are able to learn more effectively [Shawe-Taylor et al., 1998; Herbrich and Williamson, 2003].

3 Methodology

3.1 HAR Dataset

A set of trials with volunteers was required to create and develop the human activity recognition dataset. In total, 30 people with ages from 19 to 48 years participated in this research and performed a set of motion sequences comprising the 6 proposed ADL (*standing, sitting, laying, walking, walking upstairs* and *walking downstairs*). Each subject performed the experiment protocol twice, and each activity was at least performed two times on each trial to simulate repeatability (refer to Table 1 for further details). Also, a timeout of 5 seconds

No.	Static	Time (sec)	No.	Dynamic	Time (sec)
0	Start (Standing Pos)	0	7	Walk (1)	15
1	Stand (1)	15	8	Walk (2)	15
2	Sit (1)	15	9	Walk Downstairs (1)	12
3	Stand (2)	15	10	Walk Upstairs (2)	12
4	Lay Down (1)	15	11	Walk Downstairs (1)	12
5	Sit (2)	15	12	Walk Upstairs (2)	12
6	Lay Down (2)	15	13	Walk Downstairs (3)	12
			14	Walk Upstairs (3)	12
			15	Stop	0
				Total	**192**

Table 1: Protocol of activities for the HAR Experiment.

(a) (b)

Figure 1: Samsung Galaxy S2 Smartphone. (a) arrows show the axis orientation of the accelerometer. (b) the smartphone case and belt used for the experiments

in which people remained still was arranged between each activity in order to separate each task and simplify the data labeling process.

The obtained database was partitioned into training and test sets in a proportion of 70% to 30%. The partition was randomized but assuring that no samples were from the same user in both subsets. The training data was employed for training different multiclass SVM classifiers which are described in Section 3.3. The Samsung I9100 Galaxy S II smartphone was the device utilized for the experiments with its embedded triaxial accelerometer. Acceleration signals were logged at a constant rate of 50Hz which is sufficiently fast for acquiring human body motion information [Karantonis et al., 2006]. The manual labeling process was done by selecting the videos recorded from the experiments as the ground truth and comparing them with the log files of the inertial signals. Figure 1 depicts the waist-mounted smartphone used for the experiments highlighting its casing and axis orientation of its inertial sensor. Additionally, Figure 2 shows examples of two (out of the six) performed activities during the trials.

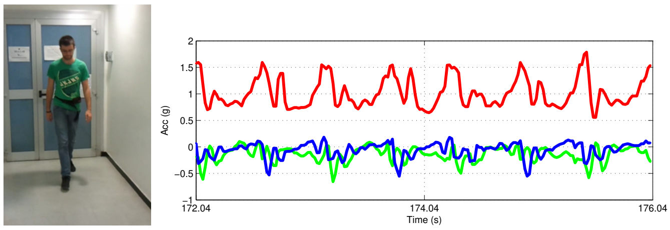

(a)

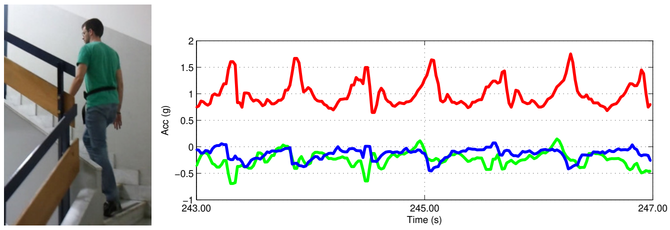

(b)

Figure 2: Human activity recognition experimental setup. Two of the six activities performed during experimentation and their triaxial acceleration signals: (a) walking, (b) walking upstairs.

3.2 Signal Processing

Sensor signals were preprocessed by the application of a series of filters for conditioning. First, noise was reduced with a median filter and a third order low-pass Butterworth filter with a cutoff frequency of 20 Hz. This frequency threshold was selected from the work presented in [Karantonis et al., 2006] which states that the energy spectrum of the human body motion lies mainly within the range of 0 Hz to 15 Hz. From these processes, a clean triaxial total acceleration $\boldsymbol{A}$ was obtained. This signal, which can be also expressed as the sum of two acceleration vectors, namely the gravitational component $\boldsymbol{G}$ and the body motion acceleration $\boldsymbol{BA}$, was segmented using another low-pass filter and assuming that the gravitational component only influences the lowest frequencies. Our experiments in segmenting these two signals found that 0.3 Hz was the optimal cutoff frequency to attain a constant gravity $\boldsymbol{G}$. This result was achieved by varying the cutoff frequency from 0.0 to 1.0 Hz in small increments of $1/40$ Hz and estimating the minimum square error of the filtered gravity signal minus the standard gravity constant $\left(9.81\ m/s^2\right)$. In addition, the acceleration time derivative $(d\boldsymbol{A}/dt)$, also known as *Jerk*, was estimated.

After segmentation, fixed-width sliding windows were captured from the preprocessed acceleration time signals, each with a span of 2.56 *sec* and an overlap of 50% which has confirmed to be successful in other HAR approaches such as in [DeVaul and Dunn, 2001; Van Laerhoven and Cakmakci, 2000]. The window length has been selected given the following reasons:

Figure 3: Activity Recognition process pipeline.

Feature Vector	
Measure	**Applied to**
SMA	*AT, JT, BAT, BJT*
Mean	AT_X, AT_Y, AT_Z
STD	AT_X, AT_Y, AT_Z
Corr	AT_X–AT_Y, AT_X–AT_Z, AT_Y–AT_Z
Entropy	AF_X, AF_Y, AF_Z, *AMF*

A:Triaxial acceleration, B:Body, J:Jerk
T:Time F:Frequency, M:Magnitude

Table 2: List of measures for computing feature vectors.

– The cadence range of an average person walking is $[90, 130]$ steps/min [Ben-Abdelkader et al., 2002] which denotes a minimum speed of 1.5 steps/sec;
– At least a full walking cycle of two steps is desirable on each window sample;
– People with slower cadence such as the disabled and elderly should also benefit from this approach. We have chosen a minimum speed of 50% the average human cadence;
– Frequency domain signals require the Fast Fourier Transform (FFT) which is optimized for power of two vectors (2.56sec $\times$ 50Hz $=$ 128cycles).

From each window, a vector of features was extracted which contained 17 features estimated from a set of measures in the time and frequency domain using previously suggested features [Bao and Intille, 2004; Lovell et al., 2007; Sama et al., 2010]: e.g. Signal Magnitude Area (*SMA*), mean, standard deviation (*STD*), entropy and signal-pair correlation (*Corr*). The Fast Fourier Transform (FFT) was used to find the frequency components for each window. The measures extracted to obtain the feature vector are depicted in Table 2. A feature vector was calculated from each experiment window sample and used as an input for the learning algorithm. The HAR process with its main components is illustrated in Figure 3.

3.3 The MultiClass HF-SVM (MC-HF-SVM) model

3.3.1 The binary HF-SVM model

Consider a dataset composed of l patterns. Each one corresponds to an ordered pair $(\boldsymbol{x}_i, y_i)$ $\forall i \in [1, ..., l]$, $\boldsymbol{x}_i \in \mathbb{R}^m$, and $y_i = \pm 1$. A standard SVM can be learned by solving a Constrained Quadratic Programming (CQP) minimization problem. This is formulated by:

$$\min_{\boldsymbol{w},b,\boldsymbol{\xi}} \quad \frac{1}{2}\|\boldsymbol{w}\|^2 + C\boldsymbol{r}^T\boldsymbol{\xi} \tag{1}$$

$$y_i\left(\boldsymbol{w}^T\boldsymbol{x}_i + b\right) \geq 1 - \xi_i, \quad \xi_i \geq 0, \quad \forall i \in [1, ..., l], \tag{2}$$

where C is the regularization parameter of the $\|\boldsymbol{w}\|^2$ term and $\boldsymbol{r}^T\boldsymbol{\xi}$ is the term representing upper bound of the number of misclassifications with $r_i = 1\ \forall i$.

This formulation is called the *primal* problem. Moreover, this problem can be reformulated and solved more easily by using the Lagrange multipliers α_i. This new representation is called the *dual* formulation and it is given by:

$$\min_{\boldsymbol{\alpha}} \quad \frac{1}{2}\boldsymbol{\alpha}^T Q\boldsymbol{\alpha} - \boldsymbol{r}^T\boldsymbol{\alpha} \tag{3}$$

$$0 \leq \alpha_i \leq C \ \ \forall i \in [1, ..., l], \quad \boldsymbol{y^T}\boldsymbol{\alpha} = 0, \tag{4}$$

where Q is the kernel matrix and is a symmetric positive semidefinite $l \times l$ matrix where $q_{ij} = y_i y_j K(\boldsymbol{x}_i, \boldsymbol{x}_j)$.

Once solved the CQP problem, new patterns can be classified by applying the SVM Feed-Forward Phase (FFP) which is given by the following formulation:

$$f(\boldsymbol{x}) = \sum_{i=1}^{l} y_i\alpha_i K(\boldsymbol{x}_i, \boldsymbol{x}) + b. \tag{5}$$

where b is the bias term and can be estimated by following the approach proposed in [Keerthi et al., 2001].

Evidently, the obtained output is invalid for its use in fixed-point arithmetic because the α_i values belong to the group of real numbers limited between 0 and C. To overcome this issue, a normalization process can be employed without affecting the classifier output sign but its magnitude and keeping the SVM accuracy unchanged. The original HF-SVM was described in [Anguita et al., 2007] and proposed the use of a new vector $\boldsymbol{\beta}$ which is defined as:

$$\beta_i = \alpha_i \frac{2^k - 1}{C}, \tag{6}$$

where k is the number of bits. Moreover the bias term b is removed as its value is quite difficult to control [Anguita et al., 2007], in order to allow a prediction with only integer parameters. This modification has no influence on the classification performance of the trained model as far as a Radial Basis Function (RBF) kernel, such as the Gaussian or the Laplacian ones, is exploited [Poggio et al., 2002]. This modification yields the following formulation:

$$\min_{\boldsymbol{\beta}} \quad \frac{1}{2}\boldsymbol{\beta}^T Q\boldsymbol{\beta} - \boldsymbol{s}^T\boldsymbol{\beta} \quad \text{s.t.} \ \ 0 \leq \beta_i \leq \frac{2^k - 1}{C} \ \ \forall i \in [1, ..., l], \tag{7}$$

where $s_i = \left(2^k - 1\right)/C \;\; \forall i \in [1, ..., l]$. Once Problem (7) is solved, $\boldsymbol{\beta}$ can straightforwardly target fixed-point arithmetic through a simple nearest-integer normalization [Anguita et al., 2007]. Note, moreover, that the cost function does not change and the second constraint of Eq. 4 is taken away from the formulation, so to set the bias term $b = 0$.

To finally have a full FFP with only integer values, it is needed to modify the representation of the the kernel $K(\cdot,\cdot)$ and the input vector $\boldsymbol{x}$ in terms of number of bits (u and v bits respectively) [Anguita et al., 2007]. This produces:

$$0 \le K(\boldsymbol{x}_i, \boldsymbol{x}) \le 1 - 2^{-u} \;\; \forall i \in [1, ..., l], \tag{8}$$

$$0 \le x_i \le 1 - 2^{-v} \;\; \forall i \in [1, ..., m]. \tag{9}$$

Consequently the modified Fixed-Point FFP formulation vector is:

$$f(\boldsymbol{x}) = \sum_{i=1}^{l} y_i \beta_i K(\boldsymbol{x}_i, \boldsymbol{x}). \tag{10}$$

In particular, we opted for a Laplacian kernel because it is more appropriate for devices with limited hardware [Anguita et al., 2007] $K(\boldsymbol{x}_i, \boldsymbol{x}_j) = 2^{-\gamma\|\boldsymbol{x}_i - \boldsymbol{x}_j\|_1}$, where the Manhattan norm is defined as $\|\boldsymbol{x}\|_1 = \sum_{i=1}^{m} |x|$ and $\gamma > 0$ is the kernel hyperparameter.

3.3.2 Generalization of HF-SVM to the Multiclass Case

The SVM FFP output range diverges on each binary classifier because these are not normalized. For extending the binary problem into a multiclass problem, we used the OVA method for comparing each class c against the other classes. However before this, a procedure to permit the comparison within the group of SVM classifiers was required. As a result, we decided to measure probability estimates for each SVM $p_c(\boldsymbol{x})$ and select the actual class c^* as the one with the highest probability output for a given test sample. The probability estimation was implemented using the approach presented in [Platt, 1999] in which the training set and the SVM model were employed to fit the FFP output values $f(\boldsymbol{x})$ with a sigmoid function of the following form:

$$p(\boldsymbol{x}) = \frac{1}{1 + e^{(\Gamma f(\boldsymbol{x}) + \Delta)}}, \tag{11}$$

in which $p(\boldsymbol{x})$ is the probability estimate, and Γ and Δ are function parameters which are properly fitted on the available learning samples.

Considering the fixed-point arithmetic limitation, the sigmoid function, which works also with real numbers, cannot be directly used for estimating $p(\boldsymbol{x})$. This can be solved by means of Look-Up-Tables (LUTs). First, a fixed number of bits t must be defined and then the probability estimates $p(\boldsymbol{x})$ can be mapped given $f(\boldsymbol{x})$ without the need of floating-point arithmetic. Our experiments have showed that $t = 8$ is suitable value for this application and it simply requires LUTs with 256 elements. The complete MC-HF-SVM process is illustrated in Figure 4.

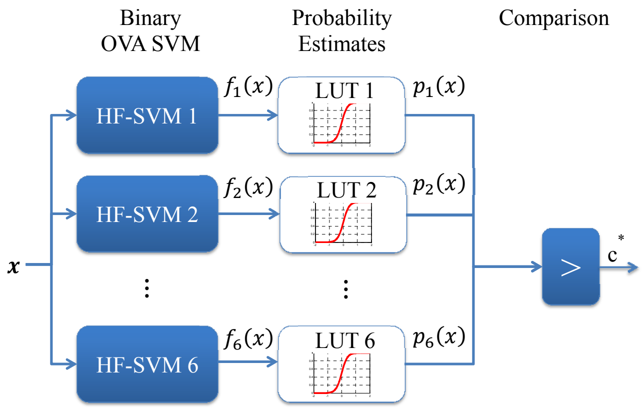

Figure 4: MC-HF-SVM using LUTs for probability estimates

3.4 HF–SVM and Statistical Learning Theory

In this section we investigate how the adoption of a fixed-point arithmetic affects the generalization ability of a classifier in the form of Eq. (10). In order to do this we describe each parameter β_i as an integer value of k bits:

$$\beta_i = \sum_{i=1}^{k-1} b_i^j 2^j, \tag{12}$$

where b_i^j is a binary valued variable $b_i^j \in \{0, 1\}$ and therefore β_i can be expressed as an integer variable such that $0 \le \beta_i \le 2^k - 1$. Since each b_i^j belongs to a finite set, for a fixed training set of cardinality l and a fixed kernel (with its hyperparameter), the number of classifiers that we can represent is finite. According to the notation of [Vapnik, 1995] we call N_f^l the number of classifiers that we can build with b_i^j, $i \in \{1, \ldots, l\}$ and $j \in \{0, \ldots, k-1\}$. Consequently we can exploit the well-known Vapnik's generalization bounds for finite hypothesis sets [Vapnik, 1995] which uses N_f^l as measure of complexity. Let then d^β be the number of nonzero parameters ($\beta_i \neq 0$) then:

$$N_f^l(k, d^\beta) \le \sum_{i=1}^{d^\beta} \binom{l}{d^\beta} \left[\left(2^k - 1\right)^{d^\beta} - \left(2^{k-1} - 1\right)^{d^\beta} \right], \tag{13}$$

where we take into account the fact that if all the parameters are even numbers, they can be divided by two without changing the class estimate. If, instead, d^b is the number of nonzero parameters, $b_i^j \neq 0$, then

$$N_f^l(k, d^b) \le \sum_{i=1}^{d^b} \binom{l\ k}{d^b}. \tag{14}$$

In the Statistical Learning Theory and Structural Risk Minimization frameworks [Vapnik, 1995], a good generalization capability on previously unseen data can be guaranteed [Vapnik, 1995; Anguita et al., 2012b] if a nested structure of the available hypothesis sets with increasing complexity is defined ($\mathcal{H}_1 \subseteq \mathcal{H}_2 \subseteq \ldots$). In this way, the generalization capability of a model can be controlled by choosing the set that achieves the best compromise between complexity and learning error.

In our case the complexity of the class can be defined through two quantities, k and d^β (or d^b). Starting from the set $\mathcal{H}_1$ with complexity $N_f^l(1,1)$ we can increase the complexity by increasing the number of bits $k \to k+1$ or by decreasing the sparsity of the representation $d^\beta, d^b \to d^\beta, d^b + 1$. In other words we have to search the best class which is as sparse as possible (smaller d^β or d^b) and represented with the minimum number of bits k. Obviously a classifier that belongs to a space with smaller complexity is also more energy efficient respect to the one that belongs to a space with higher complexity, as will also be shown in the subsequent experiments.

Increasing the complexity of the space has also direct consequence on the generalization ability of the classifier since according to the bound of Vapnik [Vapnik, 1995], which holds with probability $(1-\delta)$:

$$\pi \leq \nu + \sqrt{\frac{\ln\left[N_f^l(k,d)\right] - \ln(\delta)}{2l}} \tag{15}$$

where π is the generalization error and ν is the error obtained by the learning machine on the dataset.

This result is similar to the one presented in [Neven et al., 2008]. The important outcome of this section is that the number of bits in the HF–SVM has a strong regularization effect with an impact on the generalization ability of the classifier. Between two classifiers with approximately the same performance, we have to choose the one that can be represented with less number of bits since it is more energy efficient and it has more capacity of performing well on previously unseen data.

Finally we want to highlight that the bound in Eq. (15) is very loose since it is data independent and does not take into account the quality of the available samples for its estimation. In the last few years, proposed data dependent bounds [Bartlett and Mendelson, 2003; Bartlett et al., 2005] are becoming tighter and providing better interpretation of the generalization ability of classifiers. They have shown to work well on the performance estimation of real world problems such as in [Anguita et al., 2012b]. For these reason, the understanding of the influence of fixed–point arithmetic approaches in the estimation of these bounds is an interesting topic of research.

4 Experimental results

The performance of the MC-HF-SVM was evaluated through a collection of experiments using the HAR dataset described in this paper. The test data consisted of 789 samples made of 17 features nearly balanced with respect to each

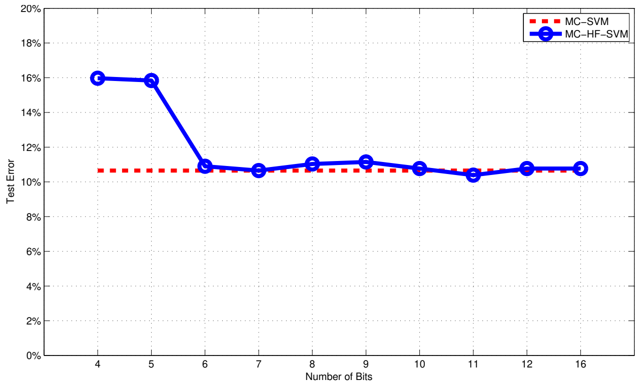

Figure 5: Comparison between error rates obtained with MC-SVM (red dotted line) and MC-HF-SVM (blue line) as k is varied.

activity. 10 MC-HF-SVM models were learned differing in the number of bits k for fixed-point representation ranging from 4 to 16 bits. Their performance was assessed in terms of test data error and compared against the standard conventional floating-point MultiClass SVM (MC-SVM). The classification results for each model are showing in Figure 5.

The test error curve shows a plateau which values appear to be stable (near 1% variation) for k ranging from 6 to 16 bits and are equivalent to the error obtained with the floating-point MC-SVM (represented with the dotted red line). In addition, the experiment shows that for this HAR dataset, $k = 6$ bits are sufficient for obtaining a recognition performance similar to the MC-SVM [see 1]. Once the numbers of bits drops below this value, the test error significantly increases by around 50%.

Moreover, it is also noticeable from the graph that some of the error values with fixed-point representation were smaller than the one found with the MC-SVM approach. This finding coincides with what observed elsewhere in literature (e.g. [Neven et al., 2008]). It is worth noting that we can remarkably reduce the number of bits (form ∞ to 8 bits) without losing the possibility of representing the functions that are characterized by good performance on the training set (as underlined in [Koltchinskii, 2006; Anguita et al., 2011, 2013]): these functions will be most likely chosen by the learning process and, then, there seem to be no reasons to search for more complex spaces. Moreover, note that few bits are required in order to represents these functions, thus contemplating an infinite-dimension space appears to be unmotivated by practical needs [Anguita et al., 2013]. Broadly speaking, the approach is theoretically feasible. In the SRM

[1] From a practical point of view, however, it is worth underlining that fixed-point libraries allow to use only values of k which are powers of 2, as will be also detailed in the forthcoming subsection: in the subsequent discussion, we will thus consider $k = 8$ as our reference value.

Method	MC-SVM						
Activity	Walking	Upstairs	Downstairs	Standing	Sitting	Laying	*Recall %*
Walking	**109**	0	5	0	0	0	*95.6*
Upstairs	1	**95**	40	0	0	0	*69.8*
Downstairs	15	9	**119**	0	0	0	*83.2*
Standing	0	5	0	**132**	5	0	*93.0*
Sitting	0	0	0	4	**108**	0	*96.4*
Laying	0	0	0	0	0	**142**	*100*
Precision %	*87.2*	*87.2*	*72.6*	*97.1*	*95.6*	*100*	***89.3***

Table 3: Confusion Matrix of the classification results on the test data using the traditional floating-point MC-SVM. Rows represent the actual class and columns the predicted class. The diagonal entries (in bold) show the number of test samples correctly classified.

Method	MC-HF-SVM $k = 8$ bits						
Activity	Walking	Upstairs	Downstairs	Standing	Sitting	Laying	*Recall %*
Walking	**109**	2	3	0	0	0	*95.6*
Upstairs	1	**98**	37	0	0	0	*72.1*
Downstairs	15	14	**114**	0	0	0	*79.7*
Standing	0	5	0	**131**	6	0	*92.2*
Sitting	0	1	0	3	**108**	0	*96.4*
Laying	0	0	0	0	0	**142**	*100*
Precision %	*87.2*	*81.7*	*74.0*	*97.8*	*94.7*	*100*	***89.0***

Table 4: Confusion Matrix of the classification results on the test data using the MC-HF-SVM with $k = 8$ bits.

framework we have to search for the simplest hypothesis space (before looking at the training set [Vapnik, 1995]) that guaranties the best trade off between accuracy on the training set and complexity of the space. Then the introduction of a bit–based hypothesis space is also encouraged by the basic ML idea to search for the simplest class of functions capable of solving the problem under examination.

In Tables 3 and 4, the confusion matrices of the MC-SVM and the MC-HF-SVM with $k = 8$ bits for the test data are depicted. In them, measures of overall accuracy, recall and precision are also given and exhibit very similar values in both approaches. Small variations are noticed in the recognition accuracy of dynamic activities within the two SVM approaches such as in the *walking downstairs* and *walking upstairs* activities, which also display some misclassifications mainly to their movement similarities. Static activities on the other hand performed better, such as *laying*, for which we reported 0% classification error. Furthermore, a small misclassification overlap was found between *standing* and *sitting*, which is attributed to the waist-mounted smartphone physical location and the difficulty to discriminate between them: this is mainly due to the slight inclination difference of the phone with respect to the vertical axis when these activities are performed and largely depends on the user's body type. These classification errors could be in some way improved by incorporating new types of features or sensors into the HAR system (e.g. inclusion of gyroscopes or additional accelerometers in different body parts), while they seem to be unrelated to the type of classification approach selected (being it a fixed-point or floating-point one).

4.1 Battery Consumption

Preliminary tests were performed on the smartphone to determine the advantages of using this novel hardware-friendly approach in terms of battery consumption: we expect that avoiding the use of the floating-point unit for complex calculations can lead to energy sparing on a stand-alone device. We used a Samsung Galaxy S II smartphone equipped with a Li-Ion 1650 mAh battery with up to 610 hours of stand-by operation and the Android Gingerbread version 2.3.4 operating system. The code was written in Java for the user interface and in C for implementing the most expensive operations such as signal processing and Machine Learning algorithms more efficiently. The use of the C programming language on Android was possible thanks to the Native Development Kit (NDK) which allows embedding native code components into Android OS applications. Most of the phone services were turned off (e.g. Wi-Fi and 3G Network) and also the phone screen was switched off as this is in general the most energy consuming phone part. The idea was to isolate this process as much as possible to obtain an approximate estimation of the battery consumption of our proposed mobile app.

The accelerometer sensor was constantly reading the triaxial signal at a fixed frequency as described in Section 3.2 in a circular buffer. Every 1.28 *sec* an interruption started the activity recognition process using the last available window sample taking into account the 50% overlap between windows and their 2.56 *sec* length. The experiments were carried out following two directions: processing time and battery consumption.

A simulation of the HAR process was implemented on the smartphone with the possibility of adjusting the number representation. This could be either fixed-point or floating-point including the unsigned default data types available in C from 8 to 64 bits. They were selected because the available libraries have only power of two number representations. The time of the activity recognition process from the sensor reading to determining the SVM FFP output was measured to estimate the average prediction rate for each approach. Table 5 shows the obtained results. It is worth highlighting the large difference between the rates using the fixed-point representation instead of the floating-point and also the proportional relationship between the number of bits used and the processing time. For instance, the 32-bits integer model outperforms in speed the 32-bit float model by almost 7 times.

An additional test was carried out aimed to measure battery consumption with the floating-point and fixed-point representations. The experiment con-

Data Type	No. Bits	No.Predictions/sec
Fixed-Point Representation		
char	8	315.35
short int	16	241.54
int	32	185.00
long int	64	141.70
Floating-Point Representation		
float	32	27.04
double	64	20.68

Table 5: Estimated prediction rates on the smartphone with basic data types.

sisted of running three times the HAR smartphone application continuously and measure the time of battery discharge from a fully charged state until a minimum level of 10% was reached. We found that the average battery time using the 32-bit float model was of 89 hours and the time with the 32-bit integer was of 112 hours. This is equivalent to an increase of 25% of the battery life when the application is running alone. These results are highly dependent on the hardware and operating system used but they are showing a trend on the improvements that can be reached with this hardware-friendly approach. For obtaining a more reliable measure of the relationship between the battery savings and processing time more experimental tests with different devices and operational conditions would be required. In current scenarios, even small savings in battery consumption make a big difference in deciding whether or not to use a mobile app, such as in cases where the HAR application is required to deliver activity information to other higher-level decision applications (e.g. phone apps for maintaining a healthy lifestyle through HAR [Lane et al., 2012]), thus implying sharing system resources. In general, we aim to build a device able to operate at least during a full day so the battery recharges can occur during the night time. These results are a good indicator of the benefits that this method can offer for saving battery life and the possibility of being integrated into devices for everyday life.

5 Conclusions

In this work we presented a novel energy efficient approach for the classification of Activities of Daily Living using smartphones. It has been constructed based on a modified Support Vector Machine model that works with fixed-point arithmetic. The proposed model was supported in terms of Structural Risk Minimization principles, where simpler models are always preferred if they have (almost) equivalent ability to learn when compared to more complex approaches.

The scope of this work is to apply the current technology for ambient intelligence applications such as in remote patient monitoring and smart environments (e.g. in long term smartphone-based activity monitoring systems). Its advantages include faster processing time, and the use of less system resources which in result provide savings in energy consumption while maintaining comparable recognition performance when compared with other traditional approaches. Also the possibility of using this approach in low-cost devices (e.g. fixed-point hardware) that could eventually be used in applications such as distributed or disposable wearable sensing.

The experimental results confirmed that it is possible to substitute the standard Multiclass SVM model with more efficient fixed-point representations but further experimentation is required to evaluate the system in more realistic conditions, such as when the smartphone system shared resources are allocated for different applications. Finally, future works will also explore algorithms, able to improve adaptability to the user and to the smartphone setting on different positions: for example, fast kernel methods [Bordes et al., 2005], allowing for online learning on the device, will be analyzed, both in terms of accuracy and ability to refine the models without (remarkably) compromising energy consumption and/or computational resources usage.

Acknowledgments

This work was supported in part by the Erasmus Mundus Joint Doctorate in Interactive and Cognitive Environments, which is funded by the EACEA Agency of the European Commission under EMJD ICE FPA n 2010-0012.

References

[Allen et al., 2006] Allen, F. R., Ambikairajah, E., Lovell, N. H., Celler, B. G.: "Classification of a known sequence of motions and postures from accelerometry data using adapted gaussian mixture models"; Physiological Measurement; 27 (2006), 10, 935.

[Anguita et al., 2012a] Anguita, D., Ghio, A., Oneto, L., Parra, X., Reyes-Ortiz, J.-L.: "Human activity recognition on smartphones using a multiclass hardware-friendly support vector machine"; Ambient Assisted Living and Home Care; volume 7657; 216–223; 2012a.

[Anguita et al., 2011] Anguita, D., Ghio, A., Oneto, L., Ridella, S.: "Selecting the hypothesis space for improving the generalization ability of support vector machines"; IEEE International Joint Conference on Neural Networks; 1169–1176; 2011.

[Anguita et al., 2012b] Anguita, D., Ghio, A., Oneto, L., Ridella, S.: "In-sample and out-of-sample model selection and error estimation for support vector machines"; IEEE Transactions on Neural Networks and Learning Systems; 23 (2012b), 9, 1390–1406.

[Anguita et al., 2013] Anguita, D., Ghio, A., Oneto, L., Ridella, S.: "A learning machine with a bit-based hypothesis space"; European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning; 2013.

[Anguita et al., 2007] Anguita, D., Ghio, A., Pischiutta, S., Ridella, S.: "A hardware-friendly support vector machine for embedded automotive applications"; IEEE International Joint Conference on Neural Networks; 1360 –1364; 2007.

[Bao and Intille, 2004] Bao, L., Intille, S.: "Activity recognition from user-annotated acceleration data"; Pervasive Computing; 2004.

[Bartlett et al., 2005] Bartlett, P., Bousquet, O., Mendelson, S.: "Local rademacher complexities"; The Annals of Statistics; 33 (2005), 4, 1497–1537.

[Bartlett and Mendelson, 2003] Bartlett, P., Mendelson, S.: "Rademacher and gaussian complexities: Risk bounds and structural results"; The Journal of Machine Learning Research; 3 (2003), 463–482.

[BenAbdelkader et al., 2002] BenAbdelkader, C., Cutler, R., Davis, L.: "Stride and cadence as a biometric in automatic person identification and verification"; IEEE International Conference on Automatic Face and Gesture Recognition; 2002.

[Berchtold et al., 2010] Berchtold, M., Budde, M., Gordon, D., Schmidtke, H., Beigl, M.: "Activity recognition service for mobile phones"; International Symposium on Wearable Computers; 2010.

[Bordes et al., 2005] Bordes, A., Ertekin, S., Weston, J., Bottou, L.: "Fast kernel classifiers with online and active learning"; Journal of Machine Learning Research; 6 (2005), 1579–1619.

[Cedras and Shah, 1995] Cedras, C., Shah, M.: "Motion-based recognition a survey"; Image and Vision Computing; 13 (1995), 2, 129–155.

[Chang and Lin, 2011] Chang, C.-C., Lin, C.-J.: "LIBSVM: A library for support vector machines"; ACM Transactions on Intelligent Systems and Technology; 2 (2011), 27:1–27:27.

[Chen et al., 2012a] Chen, L., Hoey, J., Nugent, C., Cook, D., Yu, Z.: "Sensor-based activity recognition"; IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews; 42 (2012a), 6, 790–808.

[Chen et al., 2012b] Chen, L., Nugent, C., Wang, H.: "A knowledge-driven approach to activity recognition in smart homes"; IEEE Transactions on Knowledge and Data Engineering; 24 (2012b), 6, 961–974.

[Choudhury et al., 2008] Choudhury, T., Consolvo, S., Harrison, B., Hightower, J., LaMarca, A., Legrand, L., Rahimi, A., Rea, A., Bordello, G., Hemingway, B., Klasnja, P., Koscher, K., Landay, J., Lester, J., Wyatt, D., Haehnel, D.: "The mobile sensing platform: An embedded activity recognition system"; IEEE Pervasive Computing; 7 (2008), 2, 32–41.

[Cook and Das, 2007] Cook, D. J., Das, S. K.: "How smart are our environments? an updated look at the state of the art"; Pervasive and Mobile Computing; 3 (2007), 2, 53–73.

[Cook and Das, 2012] Cook, D. J., Das, S. K.: "Pervasive computing at scale: Transforming the state of the art"; Pervasive Mobile Computing; 8 (2012), 22–35.

[DeVaul and Dunn, 2001] DeVaul, R. W., Dunn, S.: "Real-time motion classification for wearable computing applications"; Technical report; MIT Media Lab (2001).

[Epitropakis et al., 2010] Epitropakis, M., Plagianakos, V., Vrahatis, M.: "Hardware-friendly higher-order neural network training using distributed evolutionary algorithms"; Applied Soft Computing; 10 (2010), 2, 398 – 408.

[Eur, 2011] Eur, 2011: "The 2012 ageing report: Underlying assumptions and projection methodologies"; European Commission. Directorate-General for Economic and Financial Affairs (2011).

[Ganapathiraju et al., 2004] Ganapathiraju, A., Hamaker, J., Picone, J.: "Applications of support vector machines to speech recognition"; IEEE Transactions on Signal Processing; 52 (2004), 8, 2348 – 2355.

[García-Vázquez et al., 2010] García-Vázquez, J. P., Rodríguez, M. D., Tentori, M. E., na, D. S., Ángel G. Andrade, Espinoza, A. N.: "An agent-based architecture for developing activity-aware systems for assisting elderly"; Journal of Universal Computer Science; 16 (2010), 12, 1500–1520.

[Genov and Cauwenberghs, 2003] Genov, R., Cauwenberghs, G.: "Kerneltron: support vector "machine" in silicon"; IEEE Transactions on Neural Networks; 14 (2003), 5, 1426 – 1434.

[He and Jin, 2009] He, Z., Jin, L.: "Activity recognition from acceleration data based on discrete consine transform and svm"; IEEE International Conference on Systems, Man and Cybernetics; 2009.

[Herbrich and Williamson, 2003] Herbrich, R., Williamson, R.: "Algorithmic luckiness"; The Journal of Machine Learning Research; 3 (2003), 175–212.

[Herrlich et al., 2011] Herrlich, S., Spieth, S., Nouna, R., Zengerle, R., Giannola, L., Pardo-Ayala, D.-E., Federico, E., Garino, P.: "Ambulatory treatment and telemonitoring of patients with parkinsons disease"; Ambient Assisted Living; 2011.

[Irick et al., 2008] Irick, K., DeBole, M., Narayanan, V., Gayasen, A.: "A hardware efficient support vector machine architecture for fpga"; International Symposium on Field-Programmable Custom Computing Machines; 2008.

[Jatoba et al., 2008] Jatoba, L. C., Grossmann, U., Kunze, C., Ottenbacher, J., Stork, W.: "Context-aware mobile health monitoring: Evaluation of different pattern recognition methods for classification of physical activity"; International Conference of the IEEE Engineering in Medicine and Biology Society; 2008.

[José et al., 2010] José, R., Rodrigues, H., Otero, N.: "Ambient intelligence: Beyond the inspiring vision"; Journal of Universal Computer Science; 16 (2010), 12, 1480–1499.

[Karantonis et al., 2006] Karantonis, D., Narayanan, M., Mathie, M., Lovell, N., Celler, B.: "Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring"; IEEE Transactions on Information Technology in Biomedicine; 10 (2006), 1, 156–167.

[Keerthi et al., 2001] Keerthi, S. S., Shevade, S. K., Bhattacharyya, C., Murthy, K. R. K.: "Improvements to platt's smo algorithm for svm classifier design"; Neural Computation; 13 (2001), 3, 637–649.

[Khan et al., 2010] Khan, A., Lee, Y.-K., Lee, S., Kim, T.-S.: "Human activity recognition via an accelerometer-enabled-smartphone using kernel discriminant analysis"; International Conference on Future Information Technology; 2010.

[Koltchinskii, 2006] Koltchinskii, V.: "Local rademacher complexities and oracle inequalities in risk minimization."; Annals of Statistics; 34 (2006), 6, 2593–2656.

[Kwapisz et al., 2011] Kwapisz, J. R., Weiss, G. M., Moore, S. A.: "Activity recognition using cell phone accelerometers"; SIGKDD Explorations Newsletter; 12 (2011), 2, 74–82.

[Lane et al., 2012] Lane, N., Mohammod, M., Lin, M., Yang, X., Lu, H., Ali, S., Doryab, A., Berke, E., Choudhury, T., Campbell, A.: "Bewell: A smartphone application to monitor, model and promote wellbeing"; IEEE International ICST Conference on Pervasive Computing Technologies for Healthcare; 2012.

[Lara and Labrador, 2012a] Lara, O., Labrador, M.: "A mobile platform for real-time human activity recognition"; IEEE Consumer Communications and Networking Conference; 667 –671; 2012a.

[Lara and Labrador, 2012b] Lara, O., Labrador, M.: "A survey on human activity recognition using wearable sensors"; IEEE Communications Surveys Tutorials; PP (2012b), 99, 1–18.

[LeCun et al., 1995] LeCun, Y., Jackel, L., Bottou, L., Brunot, A., Cortes, C., Denker, J., Drucker, H., Guyon, I., Mller, U., Sckinger, E., Simard, P., Vapnik, V.: "Comparison of learning algorithms for handwritten digit recognition"; International Conference on Artificial Neural Networks; 1995.

[Lee et al., 2003] Lee, S., Lee, S., Jung, H.: "Real-time implementation of face recognition algorithms on dsp chip"; Audio-and Video-Based Biometric Person Authentication; 2003.

[Lee and Mase, 2002] Lee, S.-W., Mase, K.: "Activity and location recognition using wearable sensors"; IEEE Pervasive Computing; 1 (2002), 3, 24–32.

[Lesser et al., 2011] Lesser, B., Mücke, M., Gansterer, W.: "Effects of reduced precision on floating-point svm classification accuracy"; Procedia Computer Science; 4 (2011), 508–517.

[Lovell et al., 2007] Lovell, N., Wang, N., Ambikairajah, E., Celler, B. G.: "Accelerometry based classification of walking patterns using time-frequency analysis"; IEEE Annual International Conference of the Engineering in Medicine and Biology Society; 2007.

[Lukowicz et al., 2004] Lukowicz, P., Ward, J., Junker, H., Stäger, M., Tröster, G., Atrash, A., Starner, T.: "Recognizing workshop activity using body worn microphones and accelerometers"; Pervasive Computing; (2004), 18–32.

[Mannini and Sabatini, 2010] Mannini, A., Sabatini, A. M.: "Machine learning methods for classifying human physical activity from on-body accelerometers"; Sensors; 10 (2010), 2, 1154–1175.

[Mantyjarvi et al., 2001] Mantyjarvi, J., Himberg, J., Seppanen, T.: "Recognizing human motion with multiple acceleration sensors"; IEEE International Conference on Systems, Man, and Cybernetics; 2001.

[Maurer et al., 2006a] Maurer, U., Smailagic, A., Siewiorek, D., Deisher, M.: "Activity recognition and monitoring using multiple sensors on different body positions"; International Workshop on Wearable and Implantable Body Sensor Networks; 2006a.

[Maurer et al., 2006b] Maurer, U., Smailagic, A., Siewiorek, D. P., Deisher, M.: "Activity recognition and monitoring using multiple sensors on different body positions";

IEEE International Workshop on Wearable and Implantable Body Sensor Networks; 4–pp; 2006b.

[Neven et al., 2009] Neven, H., Denchev, V., Rose, G., Macready, W.: "Training a large scale classifier with the quantum adiabatic algorithm"; Arxiv preprint; (2009).

[Neven et al., 2008] Neven, H., Denchev, V. S., Rose, G., Macready, W. G.: "Training a binary classifier with the quantum adiabatic algorithm"; Arxiv Preprint; (2008).

[Platt, 1999] Platt, J. C.: "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods"; Advances in Large Margin Classifiers; 1999.

[Poggio et al., 2002] Poggio, T., Mukherjee, S., Rifkin, R., Raklin, A., Verri, A.: "b"; Uncertainty in geometric computations; chapter 11, 131–141; Springer, 2002.

[Poppe, 2010] Poppe, R.: "A survey on vision-based human action recognition"; Image and vision computing; 28 (2010), 6, 976–990.

[Ravi et al., 2005] Ravi, N., D, N., Mysore, P., Littman, M. L.: "Activity recognition from accelerometer data"; Innovative Applications of Artificial Intelligence; 1541–1546; 2005.

[Rifkin and Klautau, 2004] Rifkin, R., Klautau, A.: "In defense of one-vs-all classification"; Journal of Machine Learning Research; 5 (2004), 101–141.

[Sama et al., 2010] Sama, A., Pardo-Ayala, D. E., Cabestany, J., Rodríguez-Molinero, A.: "Time series analysis of inertial-body signals for the extraction of dynamic properties from human gait"; IEEE International Joint Conference on Neural Networks; 1–5; 2010.

[Sama et al., 2012] Sama, A., Perez-Lopez, C., Romagosa, J., Rodriguez-Martin, D., Catala, A., Cabestany, J., Perez-Martinez, D., Rodriguez-Molinero, A.: "Dyskinesia and motor state detection in parkinson's disease patients with a single movement sensor"; Annual International Conference of the IEEE Engineering in Medicine and Biology Society; 1194–1197; 2012.

[Shawe-Taylor et al., 1998] Shawe-Taylor, J., Bartlett, P., Williamson, R., Anthony, M.: "Structural risk minimization over data-dependent hierarchies"; IEEE Transactions on Information Theory; 44 (1998), 5, 1926–1940.

[Silva et al., 2012] Silva, L. C. D., Morikawa, C., Petra, I. M.: "State of the art of smart homes"; Engineering Applications of Artificial Intelligence; 25 (2012), 7, 1313 – 1321.

[Stikic et al., 2011] Stikic, M., Larlus, D., Ebert, S., Schiele, B.: "Weakly supervised recognition of daily life activities with wearable sensors"; IEEE Transactions on Pattern Analysis and Machine Intelligence; 33 (2011), 12, 2521–2537.

[Turaga et al., 2008] Turaga, P., Chellappa, R., Subrahmanian, V. S., Udrea, O.: "Machine recognition of human activities: A survey"; IEEE Transactions on Circuits and Systems for Video Technology; 18 (2008), 11, 1473–1488.

[Van Laerhoven and Cakmakci, 2000] Van Laerhoven, K., Cakmakci, O.: "What shall we teach our pants?"; International Symposium on Wearable Computers; 2000.

[Vapnik, 1995] Vapnik, V. N.: The nature of statistical learning theory; Springer-Verlag New York, 1995.

[Vinh et al., 2011] Vinh, L. T., Lee, S., Le, H. X., Ngo, H. Q., Kim, H. I., Han, M., Lee, Y.-K.: "Semi-markov conditional random fields for accelerometer-based activity recognition"; Applied Intelligence; 35 (2011), 2, 226–241.

[Wawrzynek et al., 1993] Wawrzynek, J., Asanovic, K., Morgan, N., Member, S.: "The design of a neuro-microprocessor"; VLSI for Neural Networks and Artificial Intelligence; 4 (1993), 103–7.

[Wyatt et al., 2005] Wyatt, D., Philipose, M., Choudhury, T.: "Unsupervised activity recognition using automatically mined common sense"; National Conference on Artificial Intelligence; 2005.

[Yang and Yacoub, 2006] Yang, G.-Z., Yacoub, M.: Body sensor networks; volume 6; Springer, 2006.