CLASSIFICATION OF HUMAN PHYSICAL ACTIVITIES FROM ON-BODY ACCELEROMETERS
A Markov Modeling Approach

Andrea Mannini and Angelo Maria Sabatini
Arts Lab, Scuola Superiore Sant'Anna, Piazza dei Martiri della Libertà, Pisa, Italy

Keywords: Human activity classification, Statistical pattern recognition, Accelerometers, Hidden Markov Models, Human robot interaction, Machine learning.

Abstract: Several applications demanding the development of small networks of on-body sensors, such as motion sensors, are currently investigated. Accelerometers are a popular choice as motion sensors: the reason is partly in their capability of extracting information that can be used to automatically infer the physical activity the human subject is involved, beside their role in feeding estimators of biomechanical parameters. Automatic classification of human physical activities is highly attractive for pervasive computing systems, whereas contextual awareness may ease the human-machine interaction, and in biomedicine, whereas wearable sensor systems are proposed for long-term monitoring of physiological and biomechanical parameters. This paper is concerned with the machine learning algorithms needed to perform the classification task. Hidden Markov Model (HMM) classifiers are studied by contrasting them with Gaussian Mixture Model (GMM) classifiers. HMMs incorporate the statistical information available on movement dynamics into the classification process, without discarding the time history of previous outcomes, as GMMs do. In this work, rather than considering them as models for single motor activities, we apply HMMs as models suitable for sequences of chained activities. An example of the benefits of the statistical leverage by HMMs is illustrated and discussed by analyzing a dataset of accelerometer time series.

1 INTRODUCTION

Many technical applications could greatly benefit from the availability of systems that are capable of automatically classifying specific physical activities of human beings. In this paper, either static posture, e.g., standing, or dynamic motion, e.g., walking is included in the term physical activity. The sort of contextual awareness coming from this knowledge (Brézillon, 1999) may help improving the performance of healthcare monitoring devices or promoting the development of advanced human-machine interfaces. In fact, the precise activity performed by the subject helps defining the context in which further estimation can be conducted. Consider, for instance, the problem of estimating the metabolic energy expenditure of a human subject by indirect methods (Meijer et al., 1991): these methods are reported to incur severe estimation errors in the absence of any information about the particular functional task the subject is actually involved (Meijer et al., 1991 and Bouten et al., 1997). In robotics, several applications which demand some capability by the robot controller of recognizing the user's intent are, for instance, in the field of rehabilitation engineering, where smart walking support systems are developed to assist motor-impaired persons and elderly in their efforts to stand and to walk (Yu et al., 2003 and Chuy et al., 2007), or to detect gait instabilities of the user (Sabatini et al., 2002 and Hirata et al., 2008) and minimize the risk of fall (Hirata et al., 2008).

In principle, the wearable sensors needed to elicit the contextual information would be characterized by low power consumption, small size and weight, adequate metrological specifications. Micro electromechanical systems (MEMS) motion sensors appear well matched to these requirements. The methods investigated in this paper revolve around the processing of acceleration signals acquired from small nets of MEMS accelerometers affixed to selected points of the human body.

Mannini A. and Sabatini A..
CLASSIFICATION OF HUMAN PHYSICAL ACTIVITIES FROM ON-BODY ACCELEROMETERS - A Markov Modeling Approach.
DOI: 10.5220/0003151102010208
In *Proceedings of the International Conference on Bio-inspired Systems and Signal Processing* (BIOSIGNALS-2011), pages 201-208
ISBN: 978-989-8425-35-5

A major part of this paper consists of illustrating and discussing an approach to classification of human physical activities, which is based on using Hidden Markov Models (HMMs). In principle, this approach aims at exploiting the information available on the movement dynamics, namely the capability of recognizing activities performed at the current time is related to the classification outcomes provided in the past by the classifier. Accordingly, we talk about sequential classifiers, which differ from the so-called single-frame classifiers, in the sense that the latter ones are interested to single activity primitives, in other words elementary activities are studied in isolation from the history of previously detected activities (Allen et al., 2006; Bao & Intille, 2004; Begg & Kamruzzaman, 2005; Foerster et al., 1999; Karantonis et al., 2006; Mathie et al., 2004; Ravi, 2005 and Van Laerhoven & Cakmakci, 2000).

Nowadays HMMs find applications in a large number of recognition problems, including, but not limited to, speech recognition (Rabiner, 1989), hand gesture and sign language recognition (Liang & Ouhyoung, 1995), controlling robotic tools by hand gesture (Yang et al., 1997). Concerning the human activity recognition, most studies on the application of HMMs (Babu, 2002; Martìnez-Contreras, 2009) are based on camera recordings, as shown by Yamato (1992). These studies focus on the validation of statistical models of each considered activity. In a different way, our approach is based on using lightweight wearable sensors and is oriented to exploit HMMs at a higher level. In particular, their use can be oriented towards modelling time relations between elements of a sequence of activities. Few applications of HMMs are reported in the literature as for the problem of classifying human physical activities from inertial sensors, probably because HMMs are known potentially plagued by severe difficulties of parameter estimation. In this paper we propose a way of alleviating this difficulty by adopting a supervised approach to classifier training. This approach is feasible when the data available in the training set are annotated.

2 MATERIALS AND METHODS

2.1 Datasets for Physical Activity Classification

The present work is based on analyzing the dataset of acceleration waveforms published by Bao & Intille (2004), and disclosed to us by the authors. Acceleration data, sampled at 76.25 Hz, are acquired from five bi-axial accelerometers, located at the hip, wrist, arm, ankle, and thigh. The original protocol is based on testing 20 subjects, who are requested to perform 20 activities. In this paper, we select the seven activities shown in Figure 1, giving rise to a reduced dataset, henceforth called *seven-activity* dataset. These activities involve primarily the use of the lower limbs; the rationale for their inclusion is consistent with the most important item in our current research agenda, namely the development of a system for pedestrian navigation and gait parameter estimation.

Since the research goal in Bao & Intille, (2004) is exclusively to test single-frame classifiers, the available data for each subject concern acceleration time series that are known to correspond to each activity primitive. Simulating a composite activity by a single subject in our study (virtual experiment) requires that one data frame is associated to each state of the model. The associated data frame is randomly sampled (with replacement) from the maximum number N of frames available in the reduced dataset for each primitive and subject ($18 \leq N \leq 58$). We assume that a sequence of elementary activities, say, an activity at the *motor sentence* level, can be modeled as a first-order Markov chain, composed of a finite number Q of states S_i; each state accounts for an activity primitive, say, an elementary activity at the *motor word* level. The time evolution of a first-order Markov chain is governed by the vector $\boldsymbol{\pi}$ of prior probabilities, and the transition probability matrix (TPM) $\mathbf{A}$. We opt for a subject-specific training, i.e. a distinct classifier is trained for each individual subject and we build a Q-state model ($\boldsymbol{\pi}$, $\mathbf{A}$), so as to generate motor sentences from the vocabulary of motor words shown in Figure 1 ($Q = 7$).

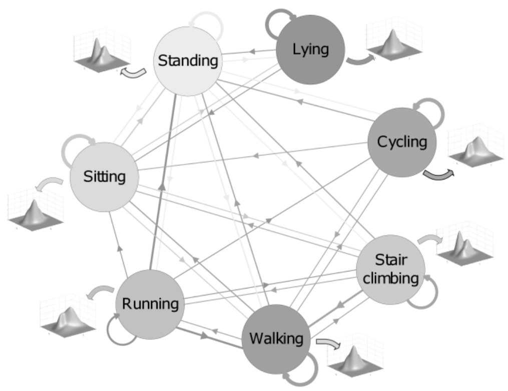

Figure 1: Scheme of a sequential classification based on HMMs.

A number $S = 20$ of virtual experiments is synthesized, each of which composed of $T = 300$ data frames. A subset of P virtual experiments is included in the training set. The procedure of synthesizing virtual experiments in the manner described above implies the existence of clear-cut borders between data frames associated to different primitives, without unknown transients between consecutive classifiable frames. This problem is managed by manual data cropping in creating the original dataset (Bao & Intille, 2004). Of course, real-life composite activities would be more complex and fuzzy, especially as for the postural transitions between different activities. In the attempt to get a more realistic picture of the HMM-based sequential classifier performance, data frames from the original dataset not included in the reduced dataset are randomly interspersed in the tested data sequences generated by the OMM, in variable proportions, from null to 1:3 (max.). The resulting garbage is managed in our system by the spurious rejection algorithm described in Section 2.5.

At the time being, the wearable system ActiNav is moving its first steps in our lab for applications in the field of pedestrian navigation and smart estimation of biomechanical parameters and it is therefore a welcome addition to our opportunities to test the developed algorithm. ActiNav revolves around an ARMadeus Board (APF27). It is equipped with an ARM9 based Freescale processor, having 128 MB of RAM, 256 MB of FLASH memory, and a 200K-gates Xilinx FPGA. A custom printed circuit board allows arming the APF27 with a 12-bit Successive Approximation Register ADC (AD7490, Analog Devices, Inc.). This converter operates up to 1 MSPS; moreover, since it is endowed with 16 analog channels, up-to-five tri-axis analog accelerometers or gyros can be integrated in ActiNav. The system with the main board (100×84×16 mm) and different sensors connected is shown in Figure 2. For the aim of this work a single tri-axis accelerometer (ADXL325, Analog Devices, Inc.) with FS = ±5 *g* is fastened on the right thigh of a single tested subject. The acquired dataset is limited to 20 sit-stand-walk sequences. This low-complexity dataset allows us in testing the proposed methods on a real sequential dataset that includes a postural transition and the incipient locomotion situation. These aspects are particularly relevant for our studies in robotic walking aids for rehabilitation.

Accelerometer data, acquired at a sampling frequency of 250 Hz, are labeled using the activity class reported by the experimenter (supervised approach). Henceforth, we refer to this reduced-complexity dataset as the *sit-stand-walk* dataset, so as to differentiate it from the *seven-activity* dataset.

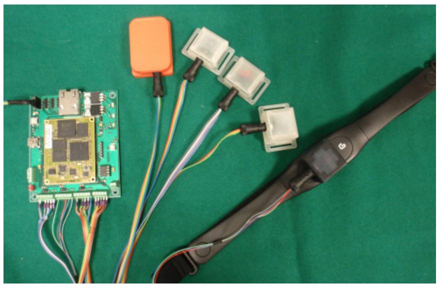

Figure 2: The ActiNav board is shown with several sensors connected to its input ports.

2.2 Data Processing: Feature Vectors

The automatic classification of acceleration data requires a pre-processing phase in which feature variables with high information content are extracted from the raw sensor data. The feature vectors are computed from acceleration samples within sliding windows with finite and constant width, henceforth called data frames.

According to the indications reported in previous works (Bao & Intille, 2004; Ravi et al., 2005), the following feature variables are selected in this paper:

- *DC component*. This feature is helpful in discriminating static postures; it is evaluated by averaging the raw samples in each data frame. One feature per data channel is obtained.
- *Energy*. This feature is helpful in assessing the motor act strength. It is evaluated as the sum of squared spectrogram coefficients within each data frame. The first coefficient that includes information about the DC component is excluded from the sum. One feature per data channel is obtained.
- *Entropy of spectrogram coefficients*. This feature is helpful in discriminating primitives that differ in frequency domain complexity (Bao & Intille, 2004). A kernel density estimator is applied to spectrogram coefficients for its determination. One feature per data channel is obtained.
- *Correlation coefficients between pair of accelerometer signals*. They are obtained by computing the dot product of pairs of frame vectors, normalized to their length, and are helpful in discriminating activities that involve motions of several body parts. A total of 55 coefficients can be computed in our application.

Before applying the classification algorithm, the computed feature vectors are selected in order to reduce the dimensionality of the problem. This is required to limit the risk of bad parameter estimation (Jain et al., 2000). In particular, we use the Pudil's algorithm that is a sequential forward-backward floating search (SFFS-SFBS) (Pudil, 1994); this algorithm uses the Euclidean distances between each pair of feature vectors of the same class in the training set as a criterion for selection.

For the *sit-stand-walk* dataset we limit ourselves to computing the DC components and the correlation coefficients. Moreover, rather than applying the Pudil's feature selection approach, we prefer to apply a feature extraction method (Jain et al., 2000). Hence, a *Principal Component Analysis* (PCA) is applied, in order to reduce the dimensionality from nine (3 DC components + 6 Correlation coefficients) to three.

2.3 Single-frame Classification

Although several single-frame classifiers can be proposed, we consider here a particular technique for single-frame classification, namely the Gaussian Mixture Model (GMM) classifier. This approach is reported by Allen et al. (2006) to achieve very promising results. In particular, the authors discuss the high adaptability of the classifier, a good feature to analyze data from subjects that are not included in the training set.

Of course, other methods for single-frame classification of human physical activity can be chosen, and they may also outperform GMMs (Mannini & Sabatini, 2010). Here, the GMM classifier is selected as the single-frame classifier of reference, in particular for its resemblance to the structure of a cHMM. As a matter of fact, the probability density of emissions of each state in a cHMM is modeled as a Gaussian mixture.

The GMM classifier first performs a parametric estimation of class-conditional probability density functions $p(x|w_i)$, which assign the probabilities of the feature vector x given its membership to the class w_i. During the training phase of a GMM classifier, class-conditional probabilities are estimated on the feature-space as Gaussian mixtures. Each feature vector x is then classified in the class yielding the highest value of $p(x|w_i)$.

2.4 cHMM-based Classification

In modeling sequences of human activities as first-order Markov chains we propose that the prior and transition probabilities that are associated to the model are empirically determined by observing the subject behavior. If the TPM and the state at the current time are known, then the most likely state that will follow is probabilistically determined. However, each activity primitive can only be observed through a set of raw sensor signals (the measured time series from on-body accelerometers, in the present case). In other terms, the states are hidden and only a second-level process is actually observable (emissions). The statistical model including the pair ($\boldsymbol{\pi}$, $\mathbf{A}$) and the emission process is an HMM. We opt for a continuous emissions approach (*continuous emissions densities* HMM, aka cHMM, Rabiner, 1989). The most common approach to the problem of modeling continuous emissions is parametric. In particular we consider mixtures of M multivariate normal distributions $N(\boldsymbol{\mu}_{jm},\boldsymbol{\Sigma}_{jm})$ that are specified by assigning the mean value vectors $\boldsymbol{\mu}_{jm}$, the covariance matrices $\boldsymbol{\Sigma}_{jm}$, and the mixing parameters matrix $\mathbf{C}$. The mixture is used to model the emissions from each state in the chain. An excellent reference source for HMMs and algorithms for their learning and testing in a recognition problem is in Rabiner (1989).

For our particular problem we consider a Q-state cHMM as represented in Figure 1 ($Q = 7$). One of the main problems may be in the high number of parameters to be identified. In fact, a Gaussian cHMM trained in a d-dimensional feature space, with Q primitives to be classified and M components for each mixture requires the specification of the following parameters:

- $\boldsymbol{\pi}$, prior probability vector, $1 \times Q$;
- $\mathbf{A}$, transition probability matrix $Q \times Q$;
- $\boldsymbol{\mu}$, set of mean value matrices, $Q \times M \times d$;
- $\boldsymbol{\Sigma}$, set of covariance matrices, $Q \times M \times d \times d$;
- $\mathbf{C}$, set of mixing parameters, $Q \times M$.

The approach to deal with the parameter identification problem is to split the training phase into two different steps: a first-level supervised training phase is followed by a second-level training phase, which is performed by running the Baum-Welch algorithm (Rabiner, 1989). Indeed, the particular problem we are facing with is typically supervised. It is also known that an inaccurate initialization of parameters could lead to suboptimal results by using the Baum-Welch algorithm, due to the presence of many local maxima in the optimization surface (Rabiner, 1989). Accordingly, the first level supervised training becomes the proposed particular way for achieving a good initialization of parameters entering the second "traditional" phase.

In order to simplify the estimation process, the parameter set is divided into two main groups, namely transition parameters (**π**, **A**) and emission parameters (**μ**, **Σ**, **C**). This separation allows us to train separately two parameter sets with reduced size, yielding a relevant reduction of the overall size of the training set. Transition parameters can be estimated through an OMM. In fact, under supervised conditions, activity labels from training set sequences, which in our model correspond to hidden states, are actually known. Emission parameters can be estimated by running a GMM classifier. The training process at the second level exploits the values of the parameters estimated during the training process at the first level, as initial values for running the Baum-Welch algorithm.

In Figure 3 a conceptual scheme of the whole sequential classification algorithm is depicted: thin lines refers to parameters, bold lines represent data frames.

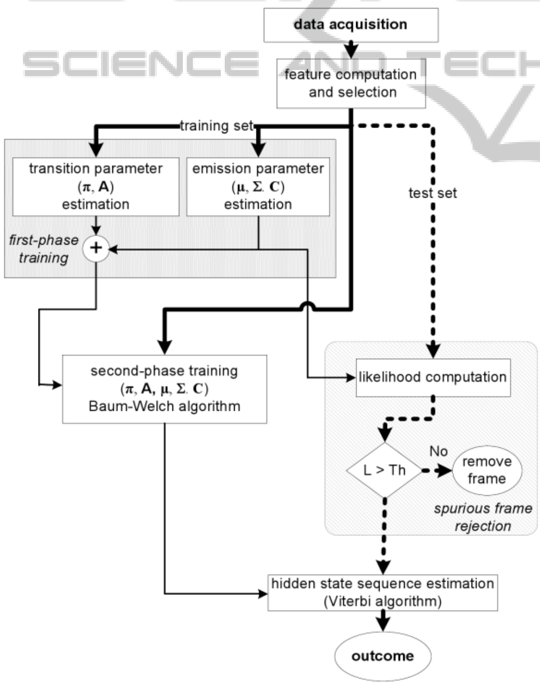

Figure 3: Block diagram of the developed cHMM-based sequential classifier.

2.5 Spurious Data Rejection

The introduced classification strategy allows us to define a criterion for automatic rejection of spurious feature vectors. If a threshold-based detector is applied to estimated class-conditional probabilities $p(x|w_i)$, it is straightforward to reject those feature vectors the classification of which is believed too uncertain, without introducing an additional model for unknown data. In fact the probability $p(x|w_i)$ in the cHMM refers to the probability of the feature vector x of being the emission of the model state w_i. If, for any feature vector, the probabilities relative to each state are below the threshold, the feature vector itself can be marked as spurious and removed, without affecting the cHMM operation. Low values of $p(x|w_i)$ are typical when unknown activities are hidden in the data presented to the classifier, or when too much uncertainty affects them.

The threshold value can be optimized upon assessment of the ROC curves; in Figure 4, the specificity-sensibility curve, averaged over the 20 subjects, is reported for the data in the *seven-activity* dataset. The threshold is settled in our application by retaining the value when the sensibility of rejection is slightly greater than the specificity.

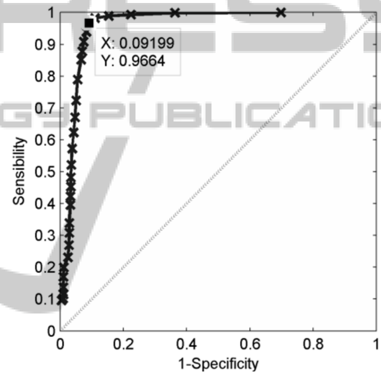

Figure 4: ROC curve obtained for different threshold values.

3 RESULTS

3.1 The *seven-activity* Dataset

After applying the Pudil's feature selection algorithm to data, the number of features is reduced from 85 to 17, namely 4 DC components of accelerations and 13 correlation coefficients are retained. The training set for the single-frame classifier is composed of K frames per class and per subject. According to the results of some preliminary testing, K = 7 turns out to be a convenient choice. Testing is performed using the remaining N–K frames available for each subject.

The number of Gaussian components of the mixture is taken M = 1, either in the GMM or the cHMM-based classifiers. Indeed the experimental evidence is in strong support of the assumption of unimodal data distributions. Algorithm testing up to

$M = 5$ indicates only marginal improvements over the simpler choice $M = 1$ discussed in the following. We consider the value $K = 7$ for the cHMM-based classifier too. The effect of the number P of motor sentences in the training set is then analyzed and results are shown in Figure 5, either in the case that the second-level training is performed or not, yielding $P = 5$ as a reasonable value for sizing the training set. A subject-specific training and test is performed for both GMM and cHMM-based classifiers. In Table 1 the classification accuracy is reported for each tested subject.

As far as the algorithm for spurious data rejection is concerned, the threshold is fixed so as to achieve sensibility Se = 96.6% and specificity Sp = 90.8%. In Table 2 the classification accuracy in the presence of spurious data and after their automatic rejection is presented.

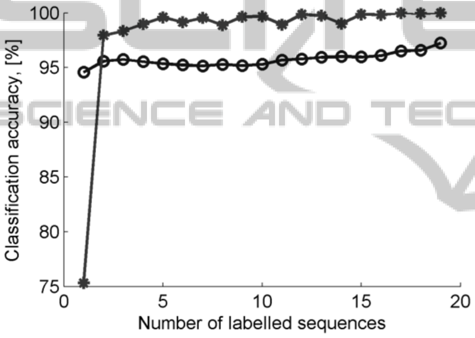

Figure 5: Classification accuracy *vs*. number of P motor sentences in the training set. o: only first-level training is applied; *: first-level training is followed by second-level training.

3.2 The *sit-stand-walk* Dataset

This dataset is processed using the same sequential classification methodology as for the *seven-activity* dataset. However, the presence of a single subject dataset requires a different validation method. A *leave-one-out* approach is followed: 20 classifiers are trained, and each time a single sequence is used to validate the classifier. Classification results in terms of recognition accuracy are reported in Table 3. As far as spurious data, there is no need to add spurious data, as described before for the *seven-activity* dataset. The spurious rejection algorithm is now applied to tag data from the *sit-stand-walk* dataset, whose reliability for classification is deemed questionable. Of course, we expect to observe a higher number of tagged data where activity transitions take place. In Figure 6 the classifier outcome and the spurious rejection effect are reported.

Table 1: Recognition accuracies (percentage values) for each subject and mean accuracy value over 20 subjects (*seven-activity* dataset, before introducing spurious data).

Subject	GMM	cHMM (First level only)	cHMM (First and second level)
1	97.6	97.2	99.6
2	93.9	95.6	99.7
3	94.9	94.9	99.7
4	99.9	96.1	99.7
5	82.1	92.2	97.8
6	91.6	89.8	99.5
7	89.5	90.5	97.7
8	98.3	90.6	99.7
9	87.2	94.5	99.6
10	95.6	96.3	99.7
11	98.3	96.2	98.8
12	90.4	89.2	98.1
13	79.9	86.7	99.2
14	92.7	87.2	98.8
15	64.3	94.2	99.6
16	98.7	97.9	98.9
17	53.6	94.8	97.6
18	67.9	81.7	83.4
19	86.9	96.5	99.5
20	75.2	98.5	99.7
Mean	86.9	93.0	98.3

Table 2: Classification accuracy (mean percentage values) in the presence of spurious data, *seven-activity* dataset.

Implementation	Classification accuracy, [%]
Without rejection of spurious data	72.1
With rejection of spurious data	95.7

Table 3: Classification accuracy (mean percentage values) after and before spurious data rejection, *sit-stand-walk* dataset.

Classifier	Classification accuracy (%)
GMM	89.7
cHMM (First level only)	86.4
cHMM (First and second level)	96.0
cHMM (With spuria rejection)	99.2

4 DISCUSSIONS AND CONCLUSIONS

Referring to the seven-activity dataset, the Pudil's feature selection scheme individuates a subset of features that simply consist of gross postural information (DC components) and movement coordination information (correlation coefficients).

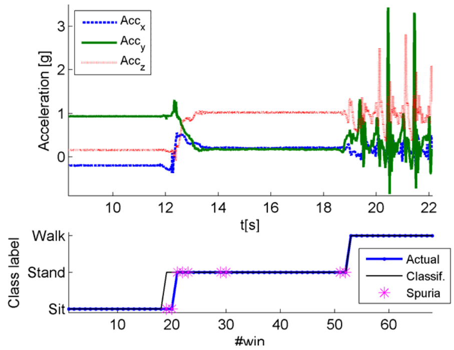

Figure 6: Classification and spurious data rejection on a sequence of the *sit-stand-walk* dataset.

Nonetheless, it is argued that energy and entropy time-domain features would be highly valuable, provided that we decide to investigate other activities, e.g., those from the set studied in (Bao & Intille, 2004) that are not considered in this paper. Our decision to concentrate on a basic vocabulary of activities is motivated by our ongoing work aimed at developing a wearable sensor system for pedestrian navigation and human locomotion rehabilitation.

The applicability of Markovian modeling to the classification of human physical activities - the subject of this paper - is demonstrated. In particular we highlight the importance of exploiting the statistical knowledge about the human motion dynamics that can be "trapped" within the Markov chain. The cHMM-based classifier, owing to the exploitation of statistical information about the activity dynamics it provides, systematically outperforms the GMM classifier (the classification accuracy, averaged across the pool of tested subjects, raises from about 87% to 98%).

A subject-specific training is considered in this work. Especially for the cHMM-based sequential classifier indeed, this approach is more appropriate than training a single net for a pool of subjects; this is because of the high variability in how humans perform a given physical activity.

The supervised training is pursued in this paper with the idea to split the process of estimating the parameters of the cHMM-based classifier into two distinct levels. We can observe that considering only the first-level training the accuracy of the cHMM-based classifier performance goes down to about 93% from 98%, still superior to the performance of the single-frame GMM classifier. Splitting the training process in two distinct levels is helpful to effectively cope with the size limitations of the training set: $P = 5$ sequences lasting each just few minutes are enough to yield a suitable training set in the present application.

A final point is related to the proposed method for managing spurious feature vectors. Most published studies, including (Bao & Intille, 2004), handle the problem of the fuzzy borders by manual data cropping. Clearly this is neither useful nor applicable if we look for a real-time system for activity classification. In our approach, the whole spurious rejection process is made automatic. When up to one third of the whole feature vectors in the data are spurious, the cHMM-based classifier accuracy is limited to about 72% in the absence of the proposed threshold-based detector. If the threshold-based detector is actually implemented the performance ramps up to about 96%.

Although being limited to three activities chained in a fixed order, and lasting few seconds only, the tests on the *sit-stand-walk* dataset show that the proposed algorithm can be applied to data in which transitions are not removed by data cropping. The beneficial effect of the dynamic information of HMMs respect to GMMs is confirmed and high classification accuracy is obtained (> 95 %). This capability encourages the application of the proposed methods even for subjects that are affected by pathologies. As it is shown in Figure 6, the spurious rejection system is able to identify those data that actually correspond to postural transitions, whose classification would be troublesome. This allows using the proposed methodology, without any particular attention to how the dataset is labeled during data acquisition.

In conclusion, a Markov modeling approach to the design of a sequential human activity classifier has been pursued. The requirements in terms of dataset size are not prohibitive, owing to the proposed subdivision of the training process into two distinct levels. The supervised machine learning algorithm also includes a very effective device for rejecting spurious feature vectors, which turns out to show high sensibility and specificity of detection.

Ongoing work will concern the extension of the proposed algorithm in the ActiNav system for applications in the field of pedestrian navigation, human robot interaction and smart estimation of biomechanical parameters.

ACKNOWLEDGEMENTS

The authors are indebted to Prof. Stephen S. Intille, for allowing them to use his acceleration dataset for the computer experiments in this paper.

REFERENCES

Allen, F. R., Ambikairajah, E., Lovell, N. H. & Celler, B. G., 2006. Classification of a known sequence of motions and postures from accelerometry data using adapted gaussian mixture models. *Physiol. Meas.*, 27, 935–951.

Babu, R. V., Anantharaman, B., Ramakrishnan, K. & Srinivasan, S., 2002. Compressed domain action classification using hmm. *Pattern Recognit. Lett.*, 23(10), 1203 – 1213.

Bao, L. & Intille, S. S., 2004. Activity recognition from user-annotated acceleration data. In Ferscha, A. & Mattern, F. *Pervasive Computing*, Springer, 301,1-17.

Bouten, C., Koekkoek, K., Verduin, M., Kodde, R. & Janssen, J., 1997. A triaxial accelerometer and portable data processing unit for the assessment of daily physical activity. *IEEE Trans. Biomed. Eng.*, 44(3):136–147.

Brézillon, P., 1999. Context in problem solving: a survey. *Knowl. Eng. Rev.*, 14(1), 47–80.

Chuy, O., Hirata, Y., Wang, Z. D. & Kosuge, K., 2007. A control approach based on passive behavior to enhance user interaction. *IEEE Trans. Rob.*, 23(5), 899–908.

Foerster, F., Smeja, M. & Fahrenberg, J., 1999. Detection of posture and motion by accelerometry: a validation study in ambulatory monitoring. *Comput. Hum. Behav.*, 15, 571 –583.

Hirata, Y., Komatsuda, S. & Kosuge, K., 2008. Fall prevention control of passive intelligent walker based on human model. *IEEE Int.Conf. on Intell. Robots Syst,* 1222–1228.

Jain, A. K., Duin, R. P. W. & Mao, J., 2000. Statistical pattern recognition: A review. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(1), 4–37.

Karantonis, D., Narayanan, M., Mathie, M., Lovell, N., & Celler, B., 2006. Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring. *IEEE Trans. Informat. Technol. Biomed.*, 10(1), 156–167.

Liang, R. & Ouhyoung, M., 1995. A real-time continuous alphabetic sign language to speech conversion VR system. *Comput. Graph. Forum*, 14(3), 67–76.

Mannini, A. & Sabatini, A. M., 2010. Machine learning methods for classifying human physical activity from on-body accelerometers. *Sensors*, 10(2), 1154–1175.

Martinez-Contreras, F., Orrite-Urunuela, C., Herrero-Jaraba, E., Ragheb, H. & Velastin, S. A., 2009. Recognizing human actions using silhouette-based hmm. *IEEE Conf. on Adv. Video Signal Based Surv.*,43–48.

Mathie, M. J., Celler, B. G., Lovell, N. H. & Coster, A. C., 2004. Classification of basic daily movements using a triaxial accelerometer. *Med. Biol. Eng. Comput.*, 42,679 – 687.

Meijer, G., Westerterp, K., Verhoeven, F., Koper, H. & Ten Hoor, F., 1991. Methods to assess physical activity with special reference to motion sensors and accelerometers. *IEEE Trans. Biomed. Eng.*, 38(3), 221–229.

Pudil, P., Novovicová, J. & Kittler, J., 1994. Floating search methods in feature selection. *Pattern Recogn. Lett.*, 15(11), 1119–1125.

Rabiner, L. R., 1989. A tutorial on hidden markov models and selected applications in speech recognition. *Proc. IEEE*, 77(2), 257–286.

Ravi, N., Dandekar, N., Mysore, P. & Littman, M. L., 2005. Activity recognition from accelerometer data. *American Association for Artificial Intelligence*, 5, 1541–1546.

Sabatini, A., Genovese, V. & Pacchierotti, E., 2002. A mobility aid for the support to walking and object transportation of people with motor impairments. *IEEE Int.Conf. on Intell. Robots Syst*, 2, 1349–1354.

Sabatini A. M., 2006, Inertial sensing in biomechanics: a survey of computational techniques bridging motion analysis and personal navigation.In Begg, R. & Palaniswami, M., *Computational Intelligence for Movement Sciences: Neural Networks and Other Emerging Techniques*;Eds.; Idea Group Pubilishing: Hershey, PA, USA, 70–100.

Van Laerhoven, K. & Cakmakci, O., 2000. What shall we teach our pants? *Proc. IEEE Int. Symposium on Wearable Computers*, 77.

Yamato J., Ohya J. & Ishii K., 1992, Recognizing human action in time-sequential images using Hidden Markov Models, *Proc. Conf. Comp. Vision and Pattern Rec.*, 379-385.

Yang, J., Xu, Y. & Chen, C., 1997. Human action learning via hidden markov model. *IEEE Trans. Syst. Man and Cyb., Part A*, 27(1), 34–44.

Yu, H., Spenko, M. & Dubowsky, S., 2003. An adaptive shared control system for an intelligent mobility aid for the elderly. *Auton. Rob.*, 15(1), 53–66.