Classification of periodic activities using the Wasserstein distance

Laurent Oudre, Jérémie Jakubowicz, Pascal Bianchi, Chantal Simon

***Abstract*—In this article we introduce a novel nonparametric classification technique based on the use of the Wasserstein distance. The proposed scheme is applied in a biomedical context for the analysis of recorded accelerometer data: the aim is to retrieve 3 types of periodic activities (walking, biking and running) from a time-frequency representation of the data. The main interest of the use of the Wasserstein distance lies in the fact that it is less sensitive to the location of the frequency peaks than to the global structure of the frequency pattern, allowing to detect activities almost independently of their speed or incline. Our system is tested on a 24-subject corpus: results show that the use of Wasserstein distance combined with some supervised learning techniques allows to compare with some more complex classification systems.**

***Index Terms*—biomedical signal processing, accelerometer signals, classification, Wasserstein distance**

I. Introduction

The estimation of physical activity (PA) can be useful in the prevention and treatment of obesity [1], [2] or aging-related disorders [3], for example. Even if there exist some reliable methods to evaluate the level of physical activity (such as oxygen uptake measurement or doubly labeled water), those are often expensive and intrusive and then do not suit for daily use. An alternative approach for the assessment of PA involves the use of unconstrained wearable systems such as accelerometers [4], [5]. The problem of PA estimation from accelerometer signals has received much attention for the latter years. Most of the times, the recognition of daily activities is performed by first dividing the accelerometer signals into frames, then calculating time- and frequency-domain features and finally labeling each frame through a classification process.

Several classification methods have been used in this context: decision trees and decision tables based on empirical or trained thresholds [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], or some more sophisticated machine learning procedures such as Artificial Neural Networks (ANN) [17], [18], [14], [19], [20], [21], k-Nearest Neighbors (k-NN) [22], [23], [24], Hidden Markov Models (HMM) and Gaussian Mixture Models (GMM) [25], [26], [27], Support Vector Machines (SVM) [28] or Naive Bayes methods [19], [29]. Some reviews or comparisons between these methods can be found in [28], [30], [31], [32]. Most of these works aim at detecting both static and dynamic activities and thus often use sensors located on the waist or on the lower back, which provide useful information on the postural orientation of the subject [11], [13]. However, several multi-sensor studies have shown that, as far as dynamic activities are concerned, higher accuracy can be achieved by placing some accelerometers on the leg [6], [7], [24], [33], [34], [35].

In this paper, we choose to focus on dynamic activities and thus to only use accelerometer data recorded on the shin. By deliberately putting aside static activities, we seek a better understanding on the frequency behavior of shin accelerometers signals during periodic activities. Also, while most papers dealing with the recognition of dynamic activities (such as walking) have investigated the features extraction process [16], [24], [33], [36], [37], we aim at designing a simple classifier which uses basic frequency-domain features but still manages to recognize the frequency patterns relative to three periodic activities : `Walking`, `Biking` and `Running`.

Our approach is as follows. First, we compute the spectrogram of the observed accelerometer data. Second, we introduce a novel classifier which allows to instantaneously relate a given vector of frequency bins to a given state/activity. Finally, we use a regularization step in order to avoid undesired erratic jumps in the estimated long-term sequence of activities. The classifier is designed by means of a learning phase, in an off-line fashion (*i.e.*, before being brought into operation). The learning phase is achieved as follows. For each possible state/activity, we construct one or several *templates*. A template is an artificial sample which can be interpreted as representative of the state. In order to select relevant templates, we introduce a method based on Nonnegative Matrix Factorization (NMF). Next, in the classification phase, each sample (*i.e.* each vector of frequency bins) is individually assigned to a state label based on its distance from the templates. We introduce a distance based on the Wasserstein metric. The main interest of the use of the Wasserstein metric lies in the fact that it is less sensitive to the location of the frequency peaks than to the global structure of the frequency pattern, allowing to detect activities almost independently of their speed or incline (see Section IV-A).

Our contribution is thus twofold.

- We introduce an original classifier based on a distance which is robust to a moderate mass transportation of the vector samples.
- We propose an original learning phase to produce relevant templates.

The paper is organized as follows. Section II introduces the experimental framework and some first data observations. Section III presents the formal framework in an generic setting and describes the general architecture of the proposed method. Section IV applies this framework in our context and introduces the Wasserstein distance that shall be used for the classification and the methods used for the selection of the templates. Section V provides experimental results and Section VI gives a discussion on the performances of our system.

Table I
SUBJECTS' CHARACTERISTICS.

Parameter	Mean ± SD	Range
Sex (M/F)	15/9	
Age (yr)	38 ± 12	19-54
Weight (kg)	74 ± 15	53.1-99.7
Height (m)	1.70 ± 0.07	1.58-1.85
Body Mass Index (kg.m^{-2})	25.4 ± 4.4	19.2-33.6

II. DATA

A. Subjects and Data Collection

Twenty-four consenting subjects were asked to perform a series of activities for 4 hours, among which walking and running on a treadmill or biking on a exercise bike. The description of the subjects' characteristics is presented on Table I. Note that among the 24 subjects, 8 are overweight (Body Mass Index (BMI) greater than 25 and strictly lower than 30) and 4 are obese (BMI greater than 30). During the whole experiment, they wore a triaxial accelerometer (MotionPod^{TM}by MOVEA) located on their shin. The device was oriented so as the x, y and z axes record respectively the medio-lateral, vertical and antero-posterior acceleration. The data acquisition was performed by the Centre de Recherche en Nutrition Humaine (Rhône-Alpes) and CEA-LETI. Raw signals were sampled at a sampling rate of 100Hz. The speeds and intensities of the different activities change according to the physical capacities of the subject so as to make sure the experiment is safe for the subject: in order to standardize the activities according to the physical capacity of the subjects, a simplified exercise tolerance test (step-test) was performed within 1 week of the series of laboratory activities. Results of these preliminary studies were used to adapt the speed of the treadmill and the resistance of the cycle ergometer in order to obtain two individually-adapted intensity levels for walking, two intensity-level for biking and one intensity-level of running. Moreover the subjects were asked to walk at the low to moderate intensity level for three inclines of the treadmill (0%, 5%, 10%).

Note that among the 24 subjects, only 16 of them did the `Running` part of the trial. For each subject:

1) The `Walking` state is composed of 4 or 5 successive walking periods on a treadmill at different speeds and inclines, each of them with an approximate duration of 5 minutes.
2) The `Biking` state is composed of 1 or 2 successive biking periods on an exercise bike at different intensities, each with a duration of 5 minutes.
3) The `Running` state (when present) corresponds to one running period (on a treadmill) of 5 minutes.
4) The `Other` state corresponds to various non-periodic activities such as staying in a static position (standing, lying, sitting...), going up and down the stairs, playing with a ball, etc., and lasts from 3 hours to 4 hours.

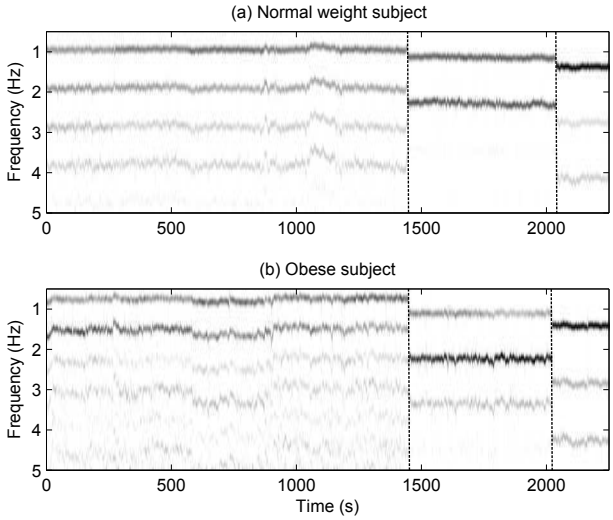

Figure 1. Example of periodic activities. The three zones correspond respectively to walking, biking and running. Normalized Short-Time Fourier Transform of the accelerometer signal recorded on the shin (anteroposterior component) for (a) a normal weight and (b) an obese subject.

B. Features extraction

The anteroposterior component of these accelerometer signals is transformed into the frequency domain, through a Short-Time Fourier Transform (STFT), calculated on 1024 samples (10.24 sec) with an overlap of 75% (which gives a new frame every 2.56 sec). We choose the anteroposterior component as it gives the clearest spectrograms (pronounced frequency structure) for the subjects of our database. Since the frequencies of most walking, running and biking movements are approximately ranged from 0.6Hz to 2.5Hz [38], we only consider the frequency bins between $f_{min} = 0.5$Hz and $f_{max} = 5$Hz. We are therefore considering $F = 92$ frequency bins.

C. First observations

By observing the spectrograms obtained on our database, we notice that the considered activities (`Walking`, `Biking` and `Running`) are characterized by a pronounced periodic structure. This claim is confirmed by Figure 1 which represents the normalized spectrogram of the raw data for the activities of interest. On this figure, the `Walking` section is in fact composed of four walking periods on a treadmill at different speeds (from 3.3 km.h^{-1} to 5.5 km.h^{-1}) and inclines (0%, 5% and 10%), and the `Biking` section is composed of two biking periods on an exercise bike at different intensities (66 W and 86 W). The spectrogram clearly shows marked spectral lines, whose locations depend on the type of activity and on the speed and incline. The existence of recognizable patterns in the spectrogram representations of these activities tends to strongly encourage the use of template-matching techniques for activity recognition.

However, in our protocol, the speed or incline of one activity varies both between subjects (according to their capacities) and within the activity. For example some subjects ran at 4 km.h^{-1} while others ran at 8 km.h^{-1}. Indeed, we observe in our data

that, even within a given state/activity, the frequency pattern is subject to significant variations according to the speed or the incline (ground, slope) or to some physical characteristics of the subject (weight, active or sedentary, etc.). In particular, the spectrograms observed for obese subjects are visibly different as those of normal weight subjects [39] (see Figure 1).

We therefore must address two important questions:

i) How to learn templates so as they best reflect the characteristics of each activity?

ii) How to build a classifier which remains robust to variations in the frequency pattern of the periodic activities?

III. Formal Framework

A. Position of the problem

Consider a finite state space $\mathcal{S}$ and an unknown sequence of T consecutive states $\boldsymbol{s} := (s_1, s_2, \cdots, s_T)$ where for each n, s_n belongs to $\mathcal{S}$. The available observations are represented as a time series $\boldsymbol{x} := (x_1, \cdots, x_T)$ where for each n, x_n belongs to an arbitrary space $\mathcal{X}$. Our aim is to estimate state sequence $\boldsymbol{s}$ from the observation $\boldsymbol{x}$.

Let assume that a database of annotated signals is available. These examples shall allow to derive some correspondence between a given state and the typical behavior of the observations for that state. It consists of a collection of Q couples of time series:

$$\left\{(\boldsymbol{x}^{(1)}, \boldsymbol{s}^{(1)}), \cdots, (\boldsymbol{x}^{(Q)}, \boldsymbol{s}^{(Q)})\right\}$$

where for any q, $(\boldsymbol{x}^{(q)}, \boldsymbol{s}^{(q)})$ consists of a known observation time series $\boldsymbol{x}^{(q)}$ and its corresponding state sequence $\boldsymbol{s}^{(q)}$. To keep notations simple we assume, without loss of generality, that each time series $\boldsymbol{x}^{(q)}$ has the same time range T. Given a new time series $\boldsymbol{x}$ associated to an *unknown* state sequence $\boldsymbol{s}$, the aim is to construct an estimate $\hat{\boldsymbol{s}}$ of $\boldsymbol{s}$, relying on the database. In this paper, we focus on estimators $\hat{\boldsymbol{s}}$ obtained by successive application of a (soft) classification step and a regularization step.

B. Classification

The key-step in the estimation of sequence $\boldsymbol{s}$ relies on the design of a sample-by-sample soft classifier. By soft classifier, we mean a mapping C of the form:

$$\begin{array}{rcl} C : \mathcal{X} & \to & \mathcal{S} \times \mathbb{R}_+ \\ x & \mapsto & (\tilde{s}(x), \tilde{\lambda}(x)) \end{array} \tag{1}$$

where $\tilde{s}(x)$ is the estimated state underlying a sample x and where $\tilde{\lambda}(x)$ is a positive scalar which stands for a level of confidence in the decision $\tilde{s}(x)$.

As described in Section I, off-the-shelf classifiers include – but are not limited to – the following techniques: k-Nearest Neighbors (k-NN), Linear Discriminant Analysis, Logistic Regression, Decision Trees, Neural Networks, Support Vector Machines (SVM) [40]. In this paper, we focus on a simple yet efficient template-based classification method. Discussion and comparison with other methods is postponed to Section V. The approach under focus goes in two steps.

1) Construction of Templates: The first step consists in the construction of a template set $\tau(s)$ for each state $s \in \mathcal{S}$. Formally $\tau(s)$ is just a subset of $\mathcal{X}^{K_s}$ where K_s is an integer which represents the number of templates available for state s. Notice that templates do *not* necessarily belong to the database. Intuitively, a template for a state s is supposed to be representative of a typical observation when the underlying state is s. There are of course many ways to build templates from the initial database. Perhaps the most immediate one is to select $\tau(s)$ as the exhaustive set of all $x_n^{(q)}$ from the database for which the underlying annotated state coincides with s. However, such an exhaustive choice for the templates leads to scan the whole database for each new observation, which is likely to be redundant, poorly efficient and time-consuming.

We defer the proposed approach for templates construction to Section IV-B.

2) Nearest Neighbor Rule: Assuming that the observation set $\mathcal{X}$ has been equipped with a well-chosen distance $d : \mathcal{X} \times \mathcal{X} \to [0, \infty)$; and the templates have already been constructed, the Voronoi tessellation is a natural choice. For any $x \in \mathcal{X}$, the output $C(x)$ of the soft classifier (1) is specified by:

$$\tilde{s}(x) = \arg\min_{s \in \mathcal{S}} \min_{y \,\in\, \tau(s)} d(x, y) \ , \tag{2}$$

and $\tilde{\lambda}(x) = \phi(\min\{d(x, y) : y \in \tau(C(x))\})$ is the corresponding confidence level, where function ϕ is a decreasing nonnegative function. The role of function ϕ is to assign high confidence levels to short distances. A standard choice is $\phi(z) = \exp(-\gamma z)$, for some fixed parameter $\gamma > 0$. Note that we shall always assume that the argument of the minimum in (2) is unique.

Of course, the performance of the estimation crucially depends on the distance d in equation (2). The choice of the distance is highly related to the nature of the observations. One of the main contributions of this paper is to introduce a relevant distance which is particularly appropriate in the case where the raw observations have a pronounced periodic behavior, as described in Section IV-A.

C. Regularization

Based on the previously described classifier C, we may now construct a sample-by-sample classification of the observation sequence $\boldsymbol{x}$. We obtain a first (temporary) estimate: $(\tilde{s}(x_1), \cdots, \tilde{s}(x_T))$. The latter estimate may be quite non regular *i.e.*, erratic jumps may occur due to possible localized classification errors. On the opposite, at least for the application of interest in this paper, the true state sequence $\boldsymbol{s}$ is supposed to be quite regular in the sense that the same state/activity usually extends over a significant period of time (in our case this period is at least equal to the size of the window i.e. 10 sec). Following [41], we propose to eventually select the final estimate $\hat{\boldsymbol{s}}$ using regularization:

$$\hat{\boldsymbol{s}} = \arg\min_{\boldsymbol{s} \in \mathcal{S}^T} \sum_{n=1}^{T} f(\tilde{\lambda}(x_n))\, \mathbf{1}_{\{\tilde{s}(x_n) \neq s_n\}} + \alpha J(\boldsymbol{s}) \ , \tag{3}$$

where $\mathbf{1}_A$ stands for the indicator function of set A, $f : \mathbb{R}_+ \to \mathbb{R}_+$ is a nondecreasing map, $\alpha > 0$ is a smoothing

parameter, and $J : \mathcal{S}^T \to \mathbb{R}_+$ is a function which penalizes non regular behaviors of state sequences. In the sequel, we use the following natural regularization term:

$$J(\boldsymbol{s}) = \sum_{n=1}^{T-1} \mathbf{1}_{s_{n+1} \neq s_n} \ .$$

Parameter α enables to control the desired weight of the regularization versus the adequacy of the final estimate to the output of the classifier. The value of α should be determined in an *ad hoc* fashion, depending on the application of interest. The factor $f(\lambda(x_n))$ in (3) tends to force the final estimate to be equal to the output of the classifier whenever the level of confidence is large. A rough but frequent choice (that we also make in the present article) is simply to set $f(\lambda) = 1$. In this case, one does not take into account the level of confidence at the output of the classifier.

IV. Classification of Periodic Activities using the Wasserstein Distance

We now return to the particular scenario of the estimation of activities from accelerometer data. The general framework of the estimation method has been described above in Section III. But in practice :

- $\mathcal{S}$ is a collection of physical activities $\mathcal{S} := \{\texttt{Walking}, \texttt{Biking}, \texttt{Running}\}$
- $x := (x_1, \cdots, x_T)$ is the squared magnitude of the Short-Time Fourier Transform (STFT). Each spectrogram frame x_n belongs to $\mathcal{X} = \mathbb{R}_+^F$ where F is the number of frequency bins considered (see II-B for details).

As pointed in Section II-C, our aim in building an efficient classifier is twofold:

- We need to choose a relevant distance d in (2) that would be specifically adapted to the application of interest and show low sensitivity to variations in speeds, inclines or subjects characteristics.
- We need to select templates which capture the frequency patterns observed in periodic activities.

A. Wasserstein Distance

The aim of the present paragraph is to select a relevant distance $d(x, y)$ between a spectrogram frame $x \in \mathcal{X}$ and a template $y \in \mathcal{X}$. Of course, the most natural choice is certainly the usual Euclidean distance. Unfortunately, the Euclidean distance is not appropriate in our case, as will be confirmed by the experimental results of Section V. An intuitive reason is the following. In our context, a sample x is a peaky function, as made clear by Figure 1: narrow peaks are located at the integer multiples of the fundamental frequency of the periodic activity. As a matter of fact, small variations in the fundamental frequency may produce large variations in the Euclidean distance between samples. As a consequence, we must rely on a different class of metrics, in such a way that a moderate shift in the fundamental frequency yields just a moderate penalization.

One of them is known as the Wasserstein distance [42]. Formally, the Wasserstein distance between two densities g and h over $[0, \pi]$ is defined by:

$$d_W(g, h) = \int_0^\pi |G(x) - H(x)| dx \tag{4}$$

where G and H denote the cumulative distribution functions associated to g and h respectively *i.e.*, $G(x) = \int_0^x g(t)dt$. Distance d_W has a nice interpretation in terms of mass transportation (see [42] for details). This propriety has been notably used in image [43] and music signal processing [44]. A simple example is given by the case where h is a shifted version of g, $h(x) = g(x - a)$. In this case $H(x) = G(x - a)$ and $d_W(g, h) = a$. In order to give further insights on d_W, we compare d_W with the Euclidean distance for two uniform densities g and h over two disjoint subintervals (a, b) and (c, d) of I with $a < b < c < d$:

$$g(t) = \frac{\mathbf{1}_{(a,b)}(t)}{b-a} \quad \text{and} \quad h(t) = \frac{\mathbf{1}_{(c,d)}(t)}{d-c}$$

On the one hand, the following holds true:

$$d_{EUC}(g, h) = \sqrt{2}$$

On the other hand, one can compare with:

$$d_W(g, h) = \frac{c+d}{2} - \frac{a+b}{2}$$

In other terms, it does not matter for the Euclidean distance how far apart are (a, b) and (c, d), their relative distance is $\sqrt{2}$; while d_W can be very small if (a, b) and (c, d) are close, or large if (a, b) and (c, d) are far-away. For instance, defining $g_n(t) = nt\mathbf{1}_{(0,1/n)}(t)$ and $h_n(t) = nt\mathbf{1}_{(1/n,2/n)}(t)$, $d_W(g_n, h_n) = 1/n$ tends to 0 when n tends to ∞, while $d_{EUC}(g_n, h_n) \equiv \sqrt{2}$.

This property of the Wasserstein distance is crucial for our application: two STFT showing the same shape but slightly shifted in frequency are considered close for the Wasserstein distance. Whereas for the Euclidean distance, or Kullback-Leibler contrast they could be considered far away.

For a practical implementation, the previous definitions have to be adapted to the discrete case. In this case, the Wasserstein distance is no longer applied to densities but to discrete probability measures, i.e. renormalized vectors. Given $x \in \mathcal{X}$ and $y \in \mathcal{X}$, by generalizing definition (4) to discrete measures, it is straightforward to show that:

$$d_W\left(\frac{x}{\|x\|_1}, \frac{y}{\|y\|_1}\right) = \sum_{f=1}^{F} \left| \sum_{j=1}^{f} \left(\frac{x(j)}{\|x\|_1} - \frac{y(j)}{\|y\|_1} \right) \right| . \tag{5}$$

However, by only using renormalized versions of x and y, we are missing the norm information, which can also be useful for classification. We therefore define an hybrid distance d for any $(x, y) \in \mathcal{X} \times \mathcal{X}$:

$$d(x, y) = d_W\left(\frac{x}{\|x\|_1}, \frac{y}{\|y\|_1}\right) + \mu \Big| \|x\|_1 - \|y\|_1 \Big| \ , \tag{6}$$

where μ is a weight parameter (for large μ, elements x and y with well distinct norm are considered far away, while for small μ their renormalized versions $x/||x||_1$ and $y/||y||_1$ only matter) and d_W is the Wasserstein distance previously defined.

B. Selection of the templates

In order to label the spectrogram frames, we propose to calculate a series of templates representative of each periodic state. Our goal is to produce a dictionary of templates whose size is large enough to capture all the characteristics of a state, but also small enough to avoid redundancy and keep the computation cost as low as possible. Two methods are used to compute these templates from annotated data: the first one is a basic computation method using averages of the spectrogram, the second one is based on a more involved dimension reduction technique (Nonnegative Matrix Factorization NMF).

1) Case A: Basic computation method: For each subject and each periodic state, we calculate the time-averaged version of all spectrogram frames belonging to the state. Note than if the number of subjects is large, this computation method also produces a large dictionary, which is likely to be redundant, noisy and to raise some practical computation issues.

2) Case B: NMF computation method: Instead of using a large dictionary which is likely to be redundant, it seems interesting to merge the available data corresponding to one state into a single and universal template. Ideally, one template per state should be sufficient to characterize an activity since the possible variations in speed, intensity or incline are supposed to be taken into account within the Wasserstein distance. Furthermore, the use of a single template makes the algorithm even less computationally demanding. The technique we use in order to learn these templates is the Nonnegative Matrix Factorization (NMF) [45]. NMF allows to decompose a matrix $\mathbf{V}$ of size $F \times T$ with nonnegative values (such as a spectrogram) into a template dictionary $\mathbf{W}$ of size $F \times D$ and an activation parameter matrix $\mathbf{H}$ of size $D \times T$ such as:

$$\mathbf{V} \approx \mathbf{WH} \tag{7}$$

where integer D is an instrumental parameter chosen so as to perform the desired dimension reduction. We refer to [45], [46] for more details on the NMF and for more precise considerations on the meaning of (7). Matrices $\mathbf{W}$ and $\mathbf{H}$ are estimated by iteratively minimizing a distance between the initial matrix $\mathbf{V}$ and the reconstruction $\mathbf{WH}$. The only constraint within the algorithm is that $\mathbf{W}$ and $\mathbf{H}$ must contain nonnegative coefficients. Many metrics have been considered in the literature (Euclidean, Kullback-Leibler, Itakura-Saito...). In particular, the Itakura-Saito divergence has provided good results when working with spectrograms [46]: this shall therefore be an appropriate metric for our task. More details on NMF algorithms are given in [46]. In our case, since the decomposition output by NMF is a priori non unique, we impose that the columns of $\mathbf{V}$ and $\mathbf{W}$ sum to 1.

In our context, we define $\mathbf{V}^{(s)}$ as the matrix containing all annotated data available for state s. More precisely, the columns of matrix $\mathbf{V}^{(s)}$ are the normalized versions $(x_n/\|x_n\|_1)$ of all spectrogram frames belonging to state s. This matrix is used as an input for the NMF algorithm, which outputs a dictionary of D templates. The templates are representatives of the frequency behavior of state s. Note that with this construction, any template has a unit norm. However, in order to be able to compute the distance (6) between a sample and a template, we must assign a norm to each template. To that end, we propose to select, among the columns of matrix $\mathbf{V}^{(s)}$, the closest frame according to the Euclidean distance, and assign the original norm of this frame to the template.

V. Results

A. Protocol

The preliminary detection of periodic states amongst the whole duration of the experiment is performed prior to the classification process by comparing the norms of the spectrogram frames $\|x_n\|_1$ to a threshold learned on the database. We assume that if the norm of a frame is lower than this threshold, the considered frame belongs to the `Other` state. This approach is motivated by the fact that all periodic states (`Walking`, `Biking`, `Running`) are supposed to have a high level of energy between $f_{min} = 0.5$Hz and $f_{max} = 5$Hz [38], contrary to non-periodic ones. Once the non-periodic frames have been detected, the remaining frames serve as input for the classification method previously described in Sections III and IV. The output sequence of estimated states (periodic and non-periodic) is then regularized through the process described in III-C.

In all the following experiments, in order to avoid overfitting, the data corresponding to the test subject is removed from the training data (24-fold cross-validation). The parameters μ (which represents the taking into account of the norm of the spectrogram frames see (6)) and α (which controls the regularization process see (3)) are estimated through a grid search with μ comprised between 10^{-6} and 10^{-4} and α between 0 and 12. The average standard deviation on all recognition scores obtained during these simulations is around 2%, which is low compared to the scores themselves (which are around 90%): this tends to show that both these parameters have a rather small impact on scores. Typical experimental values are $\alpha = 5$ and $\mu = 10^{-5}$.

B. Comparison between templates computation methods

In this section, we compare the different templates computation methods (Cases A & B) on a corpus composed of 24 subjects.

1) Case A : Basic computation method: With the basic computation method, all the frames corresponding to a state and a subject are averaged so as to produce a template. Note that in this case, the computation of the templates for the `Walking` and `Biking` states is somehow biased because they are in fact composed of different phases of the same activity with different speeds and inclines. On our corpus, this computation method outputs a dictionary of 24 templates per state *i.e.*, one for each subject (except for the `Running` state which is only performed by 16 subjects and thus yields a dictionary of only 16 templates). The obtained templates are displayed on Figure 2.

The confusion matrix obtained by using a dictionary of templates synthesized with the basic computation method is presented on Table II, along with the sensitivity (ability to avoid false negative classifications) and the specificity (ability

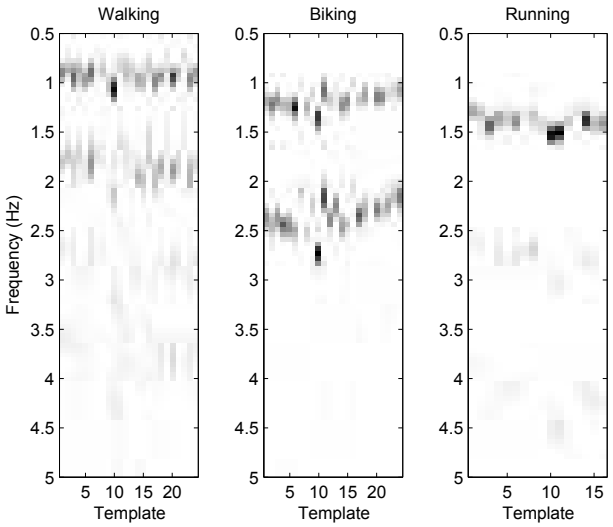

Figure 2. Case A : Templates calculated with the basic computation method (24 per periodic state except for `Running` which only contains 16 templates)

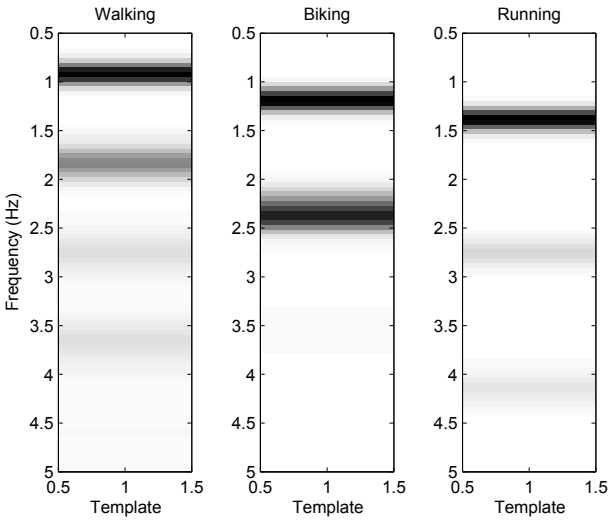

Figure 3. Case B : Templates calculated with the NMF computation method (1 template per periodic state)

Table II
CASE A. WASSERSTEIN DISTANCE. CONFUSION MATRIX OBTAINED WITH THE BASIC COMPUTATION METHOD: 1 TEMPLATE PER STATE AND PER SUBJECT.

		Detected			
		Other	**Walking**	**Biking**	**Running**
Annotated	**Other**	**94.3**	5.1	0.1	0.5
	Walking	0	**97.7**	0	2.3
	Biking	5.9	0.3	**93.7**	0
	Running	0.1	6.9	3.4	**89.6**
Sensitivity, %		94.3	97.7	93.7	89.6
Specificity, %		98.4	95.2	99.8	99.3

Table III
CASE B. WASSERSTEIN DISTANCE. CONFUSION MATRIX OBTAINED WITH THE NMF COMPUTATION METHOD: 1 TEMPLATE PER STATE.

		Detected			
		Other	**Walking**	**Biking**	**Running**
Annotated	**Other**	**94.4**	5.5	0.1	0
	Walking	0	**100**	0	0
	Biking	6.8	0	**93.2**	0
	Running	0.2	2.9	0	**96.9**
Sensitivity, %		94.4	100	93.2	96.9
Specificity, %		98.2	95.0	99.9	100

to generate true positive classifications) for each class. The scores are generally satisfying: they are approximately equal to 90% for all states.

As far as the `Other` state is concerned, we provide two interpretations for this observation. The first one is purely instrumental, and is due to slight annotations errors: there might exist slight delays between the real start/end of an activity and the corresponding annotation. These delays generate some artificial localized errors. The second possible source of confusion is due to the fact that our prior distinction between periodic and non periodic activities is somewhat arbitrary. In practice, some activities which are labeled as non periodic (such as walking up the stairs) actually show some similarities with periodic ones, and can therefore be detected as `Walking`.

The slightly lower score obtained for the `Running` state is likely to be due to the template computation method, which only uses 16 templates for this state instead of 24 for other periodic states, thus penalizing the detection of running periods. The imbalance between the number of templates is one of the weak points of this template computation method.

2) Case B: NMF computation method: In this section, the templates are learned through a supervised NMF learning process. Only one template is calculated for each state, which solves the imbalance problem described in Case A (see Section V-B1). Some examples of templates learned with NMF are provided in Figure 3. We can see on this figure that the templates calculated through NMF have naturally captured an harmonic structure, even though no such constraints have been imposed on the NMF algorithm. Also, while hybrid states such as `Walking` are composed of several periods with different fundamental frequencies, the use of NMF tends to fusion and smooth the data so as to provide only one fundamental frequency with a wider band. The results obtained on the corpus by using NMF templates are presented on Table III.

The scores obtained with this templates computation method are comparable with those obtained in Case A, except for the one obtained for the `Running` state, which is now as satisfactory as others. The fusion provided by the NMF algorithm not only allows to characterize a periodic activity with only one template, but also assesses the issue caused by the imbalance between the activities performed by the subjects. A deeper investigation shows that interestingly, 23 out of the 24 subjects have approximately a 100% detection score for all periodic states: almost all errors for these 23 subjects are caused by the `Other` state detection. The remaining subject appears to be an atypical case: the speeds and intensities of its periodic activities are clearly lower than those of other subjects. Consequently, confusions occur either in the `Other` state detection (as the norm is lower than the empirical threshold), or when taking into account of the norm through

Table IV
EUCLIDEAN DISTANCE. CONFUSION MATRIX OBTAINED WITH THE NMF COMPUTATION METHOD: 1 TEMPLATE PER STATE.

		Detected			
		Other	**Walking**	**Biking**	**Running**
Annotated	**Other**	**94.4**	0.1	5.5	0
	Walking	0	**43.5**	56.4	0
	Biking	5.8	0	**94.2**	0
	Running	0.4	0.1	16.0	**83.6**
Sensitivity, %		94.4	43.5	94.2	83.6
Specificity, %		98.4	99.0	86.1	100

parameter μ.

C. Wasserstein vs. Euclidean

In this section we aim at demonstrating that the good scores obtained on this corpus are in fact due to the Wasserstein distance's properties explained in Section IV-A. We compare the existing classification system (using the Case B templates) to its exact counterpart, only using Euclidean distance instead of Wasserstein distance. Since Euclidean distance does not impose the use of histograms, we apply it directly to the spectrogram frames (with no normalization) and we have therefore no use for parameter μ. The regularization parameter α is once again estimated through grid search. The results obtained with the Euclidean distance are displayed on Table IV and are to be compared with their Wasserstein counterpart on Table III.

According to these results, it seems clear that Wasserstein distance performs better than Euclidean distance on this task (especially for the `Walking` state). It appears that the use of Euclidean distance causes more confusions between periodic states, in particular with the `Biking` state. This can be due to the fact that the speed variations between subjects are not taken into account by the Euclidean distance: instead of searching for similar patterns, it tends to search for frames which have high amplitudes in the same frequency regions thus causing some confusions. It is particularly interesting to see that Euclidean distance fails at detecting the `Walking` state, which is the most variable in term of speeds and inclines, while Wasserstein distance is able to adapt and recognize the activity independently of the speed.

D. Comparison with classical classification systems (DT and SVM)

In this section we compare our segmentation method to two popular classification techniques: one using Support Vector Machines (SVM) and the other one using Decision Trees (DT). Most of the time, these classifiers are applied on features calculated both in the time- and frequency-domain, so as to reflect both the geometry of the posture and the physical properties of movements recorded at the back or at the waist [16], [28], [31]. However, since the geometrical properties of the shin sensor we are using in this article are arguable, it is hard to justify the use of such time- and frequency-domain features and we decide, for sake of fairness, to use the same features for all classification systems (i.e. the squared magnitude of the spectrogram).

Table V
SUPPORT VECTOR MACHINE. CONFUSION MATRIX.

		Detected			
		Other	**Walking**	**Biking**	**Running**
Annotated	**Other**	**94.1**	4.1	1.4	0.4
	Walking	0	**97.8**	2.0	0.2
	Biking	6.5	0.4	**91.0**	2.1
	Running	0.1	0.6	10.3	**89.0**
Sensitivity, %		94.1	97.8	91.0	89.0
Specificity, %		98.3	96.3	98.4	99.5

Table VI
DECISION TREE. CONFUSION MATRIX.

		Detected			
		Other	**Walking**	**Biking**	**Running**
Annotated	**Other**	**94.5**	5.4	0.1	0
	Walking	0	**100**	0	0
	Biking	6.0	1.6	**88.2**	4.2
	Running	0.1	0	0	**99.9**
Sensitivity, %		94.5	100	88.2	99.9
Specificity, %		98.4	95.0	99.9	99.8

1) Support Vector Machine: SVM-based classifiers have given good results for the classification of accelerometer signals [28], [31] (see for instance [40, chapter 12] for a nice account on SVMs). Since SVM does not properly speaking constitutes a segmentation technique but rather a classification technique, we apply the same regularization step as the one described in Section III-C. The SVM system was implemented by using the Spider Matlab package[1]. We used a standard SVM classifier using a RBF kernel ($e^{-\gamma\|x-y\|^2}$) and directly applied to the spectrogram frames. The values of the parameter of the radial basis function γ, the penalty term C and the regularization parameter α are determined through grid search. The results obtained by using the SVM system are displayed on Table V.

While the results of the SVM + regularization system are comparable with those of our Wasserstein segmentation method, they are slightly lower for all states, especially for the `Running` state. The reasons for this lower score are probably of the same kind than those explained in Section V-C: non-Wasserstein methods can only perform well when similar patterns are present in the training database and sometimes fail at adapting to different speeds or different subjects.

2) Decision Trees: Decision Tree is maybe the most common classification method used for accelerometer data [6], [9], [10], [13], [16]. We used the implementation of Matlab's Statistics Toolbox. The regularization was performed as described in Section III-C and the regularization parameter α is determined through grid search. The results obtained by using the DT system are displayed on Table VI.

We see that this method, while giving slightly lower scores for the `Biking` state, is comparable to our system. This shows

[1]http://people.kyb.tuebingen.mpg.de/spider/main.html

that our classifier, which only requires the computation of a relevant distance (see Section VI-A) can give the same results as more complex systems.

VI. Discussion

The developed system showed comparable accuracy with off-the-shelf classifiers such as SVM or DT. Yet, the originality of this method lies in the fact that it uses basic frequency-domain features and that the classification process only requires the computation of a distance. As proved by the results obtained by a similar classifier using Euclidean distance, the good results obtained by our method are mainly due to the use of the Wasserstein distance. These results show that it is possible to improve the performances of classical classifiers by adapting them to the particular task of activity recognition. In our case, the existence of frequency patterns in the spectrograms observed on the shin has allowed to build a recognition system specifically designed to classify periodic activities.

A. Complexity Analysis

A few remarks on runtime complexity are in order. According to the notations introduced in Section III & IV, computation of the Wasserstein distance has a complexity $O(F)$ where F is the number of frequency bins. It then costs $O(F(\sum_s K_s))$ to compare a new point to all existing templates. In the case where each set $\tau(s)$ consists of a singleton, the cost of a new classification is $O(F|\mathcal{S}|)$ in time. Basically, it is the best complexity one can hope for in this setting: it is linear in the problem dimension F and linear in the number of possible states $|\mathcal{S}|$. Thence it is a good candidate for a real-time implementation. The important point is that it does not depend on the size of the database.

Should we use, for instance, SVM instead, the complexity would be larger, even though SVM are considered fast in test phase. The complexity of SVM in test phase is $O(F|\mathcal{V}||\mathcal{S}|)$ where $\mathcal{V}$ is the set of support vectors, and assuming a kernel with linear cost $O(F)$ such as the polynomial or Gaussian kernels.

As seen in Section IV-B, the cost for the construction of templates in our method is linear in the size of the database. On the otherhand, the cost of SVM training is difficult to assess. Some performance-oriented solver claim a linear cost for classification training [47] but usual implementations are quadratic in the database size and linear in F.

The comparison with Decision Trees is complicated since their complexity in both the training and test phase can vary a lot.

The regularization step cost does not depend on the classifier. It can be implemented using standard dynamic programming methods. Its computational cost is $O(F|\mathcal{S}|)$.

B. Laboratory vs. real life

The main limitations to this study are due to the experimental protocol which only includes laboratory data recorded on treadmill or exercise bike. However, we consider that this study constitutes a first and compulsory step, which has allowed us to observe structured activities during at least 5 minutes. Furthermore, as pointed in Section VI-A the classification process has a low computation time. Given the latency of 10.24 seconds due to the frame effect, our method can be tested in almost real-time conditions. Also, the window length used in this paper seems reasonable for applications in daily life. Although it is impossible to predict the results that our method would obtain in real life, we can only notice that we obtained good results on a rather mixed database containing both different speeds and inclines but also different types of subject, as some of them suffered obesity. By observing detailed results obtained with Wasserstein distance and the NMF computation method (see Table VII), we notice that most of the difficulties are due to the `Biking` state which tends to give variable results, and that our system gives better results with overweight people (which is probably due to the composition of the database). Indeed, the *normal weight* category is more mixed than other categories, as it contains both sedentary and active subjects, while *overweight* and *obese* subjects are a priori sedentary. Yet, the variations between categories of subjects are reasonable and justify a future implementation for normal weight, overweight and obese people.

Table VII
Recognition scores obtained with the Wasserstein distance and the NMF computation method.

Subjects	Walking	Biking	Running
Normal weight (12)	99.9 ± 0.2	91.5 ± 27.6	95.3 ± 13.5
Overweight (8)	100 ± 0	99.8 ± 0.4	99.5 ± 0.5
Obese (4)	100 ± 0	85.0 ± 25.9	100 ± 0

Classification of the subjects according to their Body Mass Index (Normal weight: 18.5≤BMI<25, Overweight: 25≤BMI<30, Obese BMI≥30). Recognition scores are given as Mean ± SD.

VII. Conclusion

In this article, we have presented a method for segmenting and classifying periodic activities (walking, running and biking) in accelerometer signals. This method is based on a distance minimization in the frequency domain between spectrogram frames and learned templates. The novelty of our approach lies in the joint use of the Wasserstein distance, which is particularly relevant in this context, and of learning techniques (NMF) which allows to efficiently learn the characteristics of each activity. The main interest of the use of Wasserstein distance lies in its ability to recognize patterns without being too sensitive to the variations due to the specificities of the subject or to the speed of the activity. Indeed, experimental results on a 24-subject corpus show that Wasserstein distance outperforms Euclidean distance for this task, and that our classification system compares with more complex classical classification systems (SVM, DT) for the recognition of periodic activities.

Acknowledgment

This work was funded by ANR program TECSAN-SVELTE.

REFERENCES

[1] J. Levine, S. McCrady, L. Lanningham-Foster, P. Kane, R. Foster, and C. Manohar, "The role of free-living daily walking in human weight gain and obesity," *Diabetes*, vol. 57, no. 3, pp. 548–554, 2008.

[2] A. Bonomi and K. Westerterp, "Advances in physical activity monitoring and lifestyle interventions in obesity: a review," *International Journal of Obesity*, pp. 1–11, 2011.

[3] W. Zijlstra and K. Aminian, "Mobility assessment in older people: new possibilities and challenges," *European Journal of Ageing*, vol. 4, no. 1, pp. 3–12, 2007.

[4] C. Bouten, K. Koekkoek, M. Verduin, R. Kodde, and J. Janssen, "A triaxial accelerometer and portable data processing unit for the assessment of daily physical activity," *IEEE Transactions on Biomedical Engineering*, vol. 44, no. 3, pp. 136–147, 1997.

[5] K. Chen and D. Bassett Jr, "The technology of accelerometry-based activity monitors: current and future," *Medicine & Science in Sports & Exercise*, vol. 37, no. 11, pp. S490–S500, 2005.

[6] P. Veltink, H. Bussmann, W. De Vries, W. Martens, and R. Van Lummel, "Detection of static and dynamic activities using uniaxial accelerometers," *IEEE Transactions on Rehabilitation Engineering*, vol. 4, no. 4, pp. 375–385, 1996.

[7] K. Aminian, P. Robert, E. Buchser, B. Rutschmann, D. Hayoz, and M. Depairon, "Physical activity monitoring based on accelerometry: validation and comparison with video observation," *Medical and Biological Engineering and Computing*, vol. 37, no. 3, pp. 304–308, 1999.

[8] S. Lee, H. Park, S. Hong, K. Lee, and Y. Kim, "A study on the activity classification using a triaxial accelerometer," in *Proceedings of the IEEE Annual International Conference on Engineering in Medicine and Biology Society (EMBS)*, 2003, pp. 2941–2943.

[9] M. Mathie, A. Coster, N. Lovell, and B. Celler, "Detection of daily physical activities using a triaxial accelerometer," *Medical and Biological Engineering and Computing*, vol. 41, no. 3, pp. 296–301, 2003.

[10] L. Bao and S. Intille, "Activity recognition from user-annotated acceleration data," in *Pervasive Computing*, Berlin, Germany, 2004, pp. 1–17.

[11] G. Lyons, K. Culhane, D. Hilton, P. Grace, and D. Lyons, "A description of an accelerometer-based mobility monitoring technique," *Medical Engineering & Physics*, vol. 27, no. 6, pp. 497–504, 2005.

[12] Y. Ohtaki, M. Susumago, A. Suzuki, K. Sagawa, R. Nagatomi, and H. Inooka, "Automatic classification of ambulatory movements and evaluation of energy consumptions utilizing accelerometers and a barometer," *Microsystem Technologies*, vol. 11, no. 8, pp. 1034–1040, 2005.

[13] D. Karantonis, M. Narayanan, M. Mathie, N. Lovell, and B. Celler, "Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring," *IEEE Transactions on Information Technology in Biomedicine*, vol. 10, no. 1, pp. 156–167, 2006.

[14] J. Parkka, M. Ermes, P. Korpipaa, J. Mantyjarvi, J. Peltola, and I. Korhonen, "Activity classification using realistic data from wearable sensors," *IEEE Transactions on Information Technology in Biomedicine*, vol. 10, no. 1, pp. 119–128, 2006.

[15] A. Bonomi, G. Plasqui, A. Goris, and K. Westerterp, "Improving assessment of daily energy expenditure by identifying types of physical activity with a single accelerometer," *Journal of Applied Physiology*, vol. 107, no. 3, pp. 655–661, 2009.

[16] A. Bonomi, A. Goris, B. Yin, and K. Westerterp, "Detection of type, duration, and intensity of physical activity using an accelerometer," *Medicine & Science in Sports & Exercise*, vol. 41, no. 9, pp. 1770–1777, 2009.

[17] J. Baek, G. Lee, W. Park, and B. Yun, "Accelerometer signal processing for user activity detection," in *Knowledge-Based Intelligent Information and Engineering Systems*, 2004, pp. 610–617.

[18] K. Zhang, M. Sun, D. K. Lester, F. Pi-Sunyera, C. Boozer, and R. Longman, "Assessment of human locomotion by using an insole measurement system and artificial neural networks," *Journal of Biomechanics*, vol. 38, no. 11, pp. 2276–2287, 2005.

[19] K. Chang, M. Chen, and J. Canny, "Tracking free-weight exercises," *Ubiquitous Computing*, pp. 19–37, 2007.

[20] M. Ermes, J. Parkka, J. Mantyjarvi, and I. Korhonen, "Detection of daily activities and sports with wearable sensors in controlled and uncontrolled conditions," *IEEE Transactions on Information Technology in Biomedicine*, vol. 12, no. 1, pp. 20–26, 2008.

[21] J. Staudenmayer, D. Pober, S. Crouter, D. Bassett, and P. Freedson, "An artificial neural network to estimate physical activity energy expenditure and identify physical activity type from an accelerometer," *Journal of Applied Physiology*, vol. 107, no. 4, pp. 1300–1307, 2009.

[22] F. Foerster and J. Fahrenberg, "Motion pattern and posture: correctly assessed by calibrated accelerometers," *Behavior Research Methods*, vol. 32, no. 3, pp. 450–457, 2000.

[23] J. Bussmann, W. Martens, J. Tulen, F. Schasfoort, H. van den Berg-Emons, and H. Stam, "Measuring daily behavior using ambulatory accelerometry: the activity monitor," *Behavior Research Methods*, vol. 33, no. 3, pp. 349–356, 2001.

[24] S. Preece, J. Goulermas, L. Kenney, and D. Howard, "A comparison of feature extraction methods for the classification of dynamic activities from accelerometer data," *IEEE Transactions on Biomedical Engineering*, vol. 56, no. 3, pp. 871–879, 2009.

[25] F. Allen, E. Ambikairajah, N. Lovell, and B. Celler, "Classification of a known sequence of motions and postures from accelerometry data using adapted gaussian mixture models," *Physiological Measurement*, vol. 24, pp. 935–951, 2006.

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

[27] D. Pober, J. Staudenmayer, C. Raphael, and P. Freedson, "Development of novel techniques to classify physical activity mode using accelerometers," *Medicine & Science in Sports & Exercise*, vol. 38, no. 9, pp. 1626–1634, 2006.

[28] I. Gyllensten and A. Bonomi, "Identifying types of physical activity with a single accelerometer: Evaluating laboratory-trained algorithms in daily life," *IEEE Transactions on Biomedical Engineering*, vol. 58, no. 9, pp. 2656–2663, 2011.

[29] X. Long, B. Yin, and R. Aarts, "Single-accelerometer-based daily physical activity classification," in *Proceedings of the IEEE Annual International Conference on Engineering in Medicine and Biology Society (EMBS)*, 2009, pp. 6107–6110.

[30] N. Ravi, N. Dandekar, P. Mysore, and M. Littman, "Activity recognition from accelerometer data," in *Proceedings of the National Conference on Artificial Intelligence*, vol. 20, no. 3, 2005, p. 1541.

[31] S. Preece, J. Goulermas, L. Kenney, D. Howard, K. Meijer, and R. Crompton, "Activity identification using body-mounted sensors: a review of classification techniques," *Physiological Measurement*, vol. 30, no. 4, p. R1, 2009.

[32] K. Altun, B. Barshan, and O. Tunçel, "Comparative study on classifying human activities with miniature inertial and magnetic sensors," *Pattern Recognition*, vol. 43, no. 10, pp. 3605–3620, 2010.

[33] F. Foerster, M. Smeja, and J. Fahrenberg, "Detection of posture and motion by accelerometry: a validation study in ambulatory monitoring," *Computers in Human Behavior*, vol. 15, no. 5, pp. 571–583, 1999.

[34] S. Lee and K. Mase, "Activity and location recognition using wearable sensors," *IEEE Pervasive Computing*, vol. 1, no. 3, pp. 24–32, 2002.

[35] K. Zhang, P. Werner, M. Sun, F. Pi-Sunyer, and C. Boozer, "Measurement of human daily physical activity," *Obesity*, vol. 11, no. 1, pp. 33–40, 2003.

[36] M. Sekine, T. Tamura, M. Akay, T. Fujimoto, T. Togawa, and Y. Fukui, "Discrimination of walking patterns using wavelet-based fractal analysis," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 10, no. 3, pp. 188–196, 2002.

[37] N. Wang, E. Ambikairajah, N. Lovell, and B. Celler, "Accelerometry based classification of walking patterns using time-frequency analysis," in *Proceedings of the IEEE Annual International Conference on Engineering in Medicine and Biology Society (EMBS)*, Lyon, France, 2007, pp. 4899–4902.

[38] M. Henriksen, H. Lund, R. Moe-Nilssen, H. Bliddal, and B. Danneskiod-Samsře, "Test-retest reliability of trunk accelerometric gait analysis," *Gait & Posture*, vol. 19, no. 3, pp. 288–297, 2004.

[39] P. Spyropoulos, J. Pisciotta, K. Pavlou, M. Cairns, and S. Simon, "Biomechanical gait analysis in obese men," *Archives of Physical Medicine and Rehabilitation*, vol. 72, no. 13, pp. 1065–1070, 1991.

[40] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*. Springer Verlag, 2009.

[41] A. Tikhonov and V. Arsenin, *Solutions of ill-posed problems*. Washington: Winston and Sons, 1977.

[42] S. Rachev and L. Rüschendorf, *Mass transportation problems*. Springer Verlag, 1998, vol. 1.

[43] J. Rabin, G. Peyré, J. Delon, and M. Bernot, "Wasserstein barycenter and its application to texture mixing," in *Proceedings of the Third International Conference on Scale Space and Variational Methods in Computer Vision (SSMV 11)*, vol. 8, 2011.

[44] I. Katsutoshi, G. Masataka, K. Kazunori, O. Tetsuya, and G. Hiroshi, "Query-by-example music information retrieval by score-informed source separation and remixing technologies," *EURASIP Journal on Advances in Signal Processing*, vol. 2010, 2011.

[45] D. Lee and H. Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
[46] C. Févotte, N. Bertin, and J. Durrieu, “Nonnegative matrix factorization with the itakura-saito divergence: With application to music analysis,” *Neural Computation*, vol. 21, pp. 793–830, 2009.
[47] T. Joachims, “Training linear svms in linear time,” in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2006, pp. 217–226.

Laurent Oudre (M’10) was born in France in 1985. He graduated from Supélec, Gif-sur-Yvette, France in 2007 and received the M.Sc. degree in Communications and Signal Processing at Imperial College London, UK in 2007. He received his Ph.D. degree in Signal Processing at TELECOM ParisTech, Paris, France in 2010. Since February 2012, he is a post-doctoral researcher in sound processing at Ecole Normale Supérieure de Cachan, France.

His research interests focus on signal processing (sound, image and biomedical) and pattern recognition.

Jérémie Jakubowicz received the MS degree (2004) and the PhD degree (2007) in applied mathematics from the Ecole Normale Supérieure de Cachan. Since 2011 he is an Assistant Professor in the SAMOVAR lab at Télécom SudParis and an associate researcher at the CNRS.

His current research interests include distributed statistical signal processing, image processing and data mining.

Pascal Bianchi was born in 1977 in Nancy, France. He received the M.Sc. degree of Supélec-Paris XI in 2000 and the Ph.D. degree of the University of Marne-la-Vallée in 2003. From 2003 to 2009, he was an Associate Professor at the Telecommunication Department of Supélec. In 2009, he joined the Statistics and Applications group at LTCI-Telecom ParisTech.

His current research interests are in the area of statistical signal processing for sensor networks. They include decentralized detection, quantization, stochastic optimization, and applications of random matrix theory.

Chantal Simon was born in Luxembourg in 1953. She graduated from the Medical School of Strasbourg, France, in 1978 and received her Ph.D. in Human Physiology in 1998. She is actually Professor of Nutrition at the University Hospital of Lyon, France and is co-leader of a research team in the INSERM Laboratory CARMEN (Cardiovascular diseases, Metabolism and Nutrition) in Lyon, France. She is a member of the scientific board of the French National Nutrition and Health Program.

Her research interests focus on the implication of physical inactivity/activity in the physiopathology of obesity and insulin resistance, which she studies using a multidisciplinary approach combining integrative physiology, molecular biology, energetics in clinical settings and state of the art bio-logging adapted to the free living condition.