Why Your Heart Was Beating: Poster

Waltenegus Dargie
Faculty of Computer Science, Technical University of Dresden, 01062 Dresden, Germany
waltenegus.dargie@tu-dresden.de

***Abstract*—The use of wireless electrocardiograms, wearable as well as implants, enable long-term and unobtrusive monitoring of patients in their everyday living and working environments. If enriched by environmental contexts, these devices can be vital for early detection of cardiovascular diseases. Often cardiologists encourage patients to keep medical journals in order to contextualise the measurements of electrocardiograms. Experiences show, however, journal entries can be inconsistent or incomplete. In this paper we associate the measurements of a wireless electrocardiogram with the measurements of inertial sensors in order to reason about the activities of a person. We put together the raw measurements and their wavelet transform in a three-way tensor and apply tensor decomposition to uncover hidden features which can be vital for detecting the underlying activities. We model and reason about six everyday activities, namely, *cycling*, *climbing* up and down a staircase, *jumping*, *push-ups*, *running*, and *skipping*.**

***Index Terms*—Activity recognition, inertial sensors, tensor decomposition, wireless electrocardiogram**

I. Introduction

The electrocardiogram is a device which measures cardiac action potentials [1]. Under normal circumstances, the **P**, the **QRS**, and the **T** waves are produced in a single heartbeat. The former is associated with atrial contraction, the middle with ventricular contraction, and the latter with ventricular depolarisation. Understanding the occurrence, sequence, frequency, and time interval between these waves is crucial to determine various cardiac conditions.

A wide range of wireless electrocardiograms have been developed both by the research community and the industry to monitor patients in residential and rehabilitation settings [2]–[5]. The aim is to enable long-term monitoring while patients freely move and carry out everyday activities. This way, the possibility of observing symptoms which may otherwise remain hidden during clinical diagnosis increases.

Often cardiac patients are encouraged to keep medical journals in order to associate environmental, biomedical, and physiological factors – such as range of motion, pain, fatigue, headache, irritability, etc. – with ECG measurements, so that doctors can establish correlation between symptoms and potential causes and better interpret ECG measurements. Medical journalling, however, is often subjective and may contain inconsistent or incomplete reports. In this paper we fuse together the measurements of inertial sensors (3D accelerometers and 3D gyroscope) with the measurements of a wireless ECG in order to reason about underlying physical activities (or level of exertions). Our aim is to objectively explain the measurements of a wireless ECG.

This work is partially funded by the German Research Foundation (DFG) under project agreement: DA 1211/7-1 (RoReyBaN).

The remaining part of this paper is organised as follows: In Section II, we review related work. In Section III, we discuss the measurement and experiment settings as well as the preprocessing of the measurement sets. In Section IV we discuss dimensionality reduction techniques as a mechanism to efficiently model and analyse the sensed data. In Section V, we present the results of our analysis. Finally, in Section VI, we give concluding remarks.

II. Related Work

The usefulness of wireless electrocardiograms in residential and clinical settings has been the focus of research in the recent past. An overview of their association with different everyday human activities can be found in [6].

Li et al. [7] employ measurements of a 3D accelerometer and a wireless ECG to reason about motion artefacts and the activity of a user – a work similar to ours. The authors employ various time and frequency domain features as well as data fusion strategies. Thus, they combine the Hermite polynomial expansion [8], [9] and principal component analysis to model the temporal characteristics of regular cardiac activities and artefacts, respectively. Likewise, a combination of different time domain features and a support vector machine (VSM) is used to model the accelerometer measurements. In parallel, the authors extract cepstral features from ECG and accelerometer signals and combine these features using Gaussian mixture models (GMMs). Finally, the multi-modal components (ECG and accelerometer) and the multi-domain models (time domain SVM and cepstral domain GMM) are fed into a scoring function which then discriminates between different activities (lying, walking, sitting, running, fidgeting). The authors report a classification accuracy ranging from 79.3% to 97.3%.

Leutmezer et al. [10] employ a wireless ECG for detecting major epileptic seizures in real-time, relying on changes in cardiac rhythm. The system comprises of an ECG monitor and a control unit. The control unit acts as a gateway between an external clinical application and the ECG. The authors propose algorithms for detecting major epileptic seizures (tonic-clonic, generalized tonic, clonic, or hypermotor) using variations of the heart rate. The heartbeat detection algorithm relies on the continuous wavelet transformation and adaptive thresholding to detect heartbeats in ambulatory conditions even during an intense body motion. Real-time analysis of the ECG data is achieved by processing the data stream locally on the ECG monitor. Alarms are transmitted over a wireless link to the

control unit which is synchronized with a video surveillance system. The prototype is evaluated on four healthy volunteers in their home environment and on three epilepsy patients at medical facility in the Netherlands. The subjects were monitored while sleeping, for a total of 50 nights. Due to low positive predictive values in one of the patients, the authors suggested to improve their algorithm by adding other modalities, such as a 3D accelerometers, REM, and EMG sensors.

Rincon et al. [11] propose a wearable ECG that detects atrial fibrillation episodes in near real-time. Atrial fibrillation (AF) happens when disorganized electrical signals cause the heart's atria to contract very fast and irregularly resulting a desynchronization between the atria and the ventricles and causing an inefficient pump of blood [12]. Most AF symptoms are caused by a poorly controlled or irregular ventricular rate. The associated risk of stroke, dementia, heart failure, and death is high in patients who have a history of AF. Hence, early detection of AF is important to ensure immediate treatments. Since it is not known when AF can occur (it may also be asymptomatic), continuous monitoring and real-time diagnosis are crucial. AF episodes can be seen from ECG waveforms by checking for anomalies. Rincon et al.'s proposal is based on the analysis of heart rate and checking the absence of the **P** wave in the ECG waveforms. For this task, the authors propose two algorithms: one for analysing the heart rate and the other for detecting the **P** wave. The outputs of the algorithms are combined using fuzzy logic in order to classify the analysis as normal or AF. The algorithms are implemented on the Shimmer platform (ref. to Section III). The performance of the proposed system is evaluated using the MIT-BIH database [13]. The evaluation results reveal that the proposed AF detection algorithm achieves 96% sensitivity and 93% specificity.

Our approach complements proposed or existing approaches in two ways. Firstly, our measurement space offers nine degrees-of-freedom (three channels measuring ECG, three channels measuring acceleration, and three channels measuring angular velocity), thus enriching our model with evidence pertaining to human cardiology and physical activities. Secondly, we combine temporal, spectral and spatial aspects using tensors, a seamless and intuitive way of combining complementary evidence.

III. Data Acquisition and Preprocessing

We employed the Shimmer (v. 3) platform[1] for acquiring the measurement sets. The platform integrates a 5-lead wireless ECG, a 3D accelerometer, a 3D gyroscope, and a 3D magnetometer. The ECG measures cardiac action potentials at three abstract positions triangulating the heart, namely, left leg-left arm (LL-LA), left leg-right arm (LL-RA), and left arm-right arm (LA-RA). An additional measurement taken from the centre of the heart serves as a reference. So, there are four

[1]http://www.shimmersensing.com/products/ecg-development-kit (Last visited on January 08, 2019: 14:20).

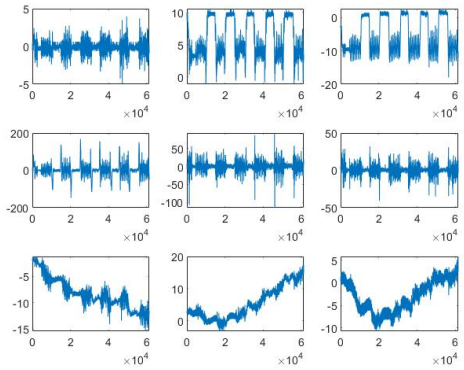

Fig. 1. The standardised measurement sets originating from a 3D accelerometer, a 3D gyroscope, and a 5-lead wireless electrocardiogram. The activity associated with the measurements was push-up.

ECG channels all together and six additional channels measuring rectilinear (3-dimensional acceleration) and curvilinear (angular velocity) motions. We abandoned the measurements of the magnetometer sensor because they were unreliable.

The Shimmer platform supports the synchronous sampling of all the channels. Its disadvantage is that the location from where the inertial measurements are taken and the location where the electrodes of the wireless electrocardiograms are placed are different, because the inertial sensors are not embedded in the electrodes. Hence, the motion affecting the individual electrodes is an approximation.

We identify six different everyday activities a subject wearing a wireless electrocardiogram may undertake – cycling, climbing up and down a staircase, jumping, doing push-ups, running, and skipping. Fig. 1 displays the row measurement sets originating from a subject performing push-ups. As can be seen, the measurement sets reveal that the subject rested between different push-up sessions. We used the intermissions to clearly mark the beginning and end of each session. Thus, we segmented the measurement sets into six push-up activity sessions. We use three of them to train our model and the other three to test our model.

In order to avoid bias arising from the heterogeneity of the row data and to give equal statistical significance to all the channels, we standardised each measurement set:

$$\mathbf{x}_n = \frac{(\mathbf{x} - \eta_x)}{\sigma_x} \tag{1}$$

where η_x is the mean and σ_x is the standard deviation of $\mathbf{x}$, respectively. Now $\mathbf{x}_n$ has a zero mean and a standard deviation of one. It should be noted that $\mathbf{x}_n$ and $\mathbf{x}$ are, statistically speaking, equivalent.

IV. Dimensionality Reduction

The research questions we wish to address in combining the measurements of the inertial sensors with the measurements of the ECG can be stated as follows:

- Do the measurements correlate well?

- Is it possible to uncover hidden features common to all the channels from the row measurements which can be useful for reasoning about the underlying activities?

We assert that these questions can be answered affirmatively and efficiently if we employ dimensionality reduction techniques. The justifications are the following:

- Since we have nine channels, the amount of data which can be generated in a short time is appreciably large.
- These data potentially exhibit strong correlations with one another, in which case dimensionality reduction techniques can reduce the overall size significantly without a significant lose of the data's statistical properties.

In the subsequent subsections, we give a brief introduction to 2D and 3D dimensionality reduction techniques – singular value decomposition and tensor decomposition –, highlighting their merit and demerits. Then, in Section V, we present in detail how we employed tensor decomposition to process the measurements of the inertial and ECG sensors to determine the activity of a person. Readers who have a sound background in dimensionality reduction techniques can skip the remainder of this section.

A. Singular-Value Decomposition

The samples taken during a single, continuous activity can be put together to form a matrix $\mathbf{X}$ of size $n \times m$, where n refers to the sample size and m refers to the number of channels. Since we are concerned with a single activity, we assume that the samples exhibit correlation in both dimensions of the matrix. Therefore, the matrix can be described by a few number of statistically independent factors. If we apply singular value decomposition (SVD) on $\mathbf{X}$, the result will be:

$$\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^{\mathbf{T}} \tag{2}$$

where $\mathbf{U}$ $(n \times r)$ and $\mathbf{V}$ $(m \times r)$ are orthonormal matrices and Σ $(r \times r)$ is a non-negative diagonal matrix whose entries $\sigma_{ii}, 1 \leq i \leq r$ are sorted in decreasing order. The diagonal entries of Σ are called the singular values of $\mathbf{X}$ and encode the distinct features of $\mathbf{X}$ and their relative significance. The advantage of using SVD is that:

1) One does not need to make any assumption as regards the hidden features. Their number and significance is automatically revealed by the decomposition process in the diagonal matrix.
2) One can express the original matrix $\mathbf{X}$ as the summation of many matrices as follows:

$$\mathbf{X} = \sum_{r=1}^{R} \sigma_{rr}\mathbf{u}_r \circ \mathbf{v}_r \tag{3}$$

where $\circ$ indicates the vector (outer) product and $\mathbf{u}_r$ and $\mathbf{v}_r$ refer to the r-th column of the matrices $\mathbf{U}$ and $\mathbf{V}$, respectively. Note that the relevance of each component is determined by the relevance of σ_{ii}.

3) If the samples of the original matrix exhibit strong correlations, then, $\mathbf{X}$ can be approximated by taking the first K components only:

$$\mathbf{X} \approx \sum_{r=1}^{K} \sigma_{rr}\mathbf{u}_r \circ \mathbf{v}_r \tag{4}$$

for $K < R$. If the difference in magnitude between the successive σ_{rr} entries is significantly large, then a strong correlation is identified in the original matrix and, hence, the error resulting from our approximation will be significantly small.

One of the limitations of working with SVD is that it is only two dimensional. Since one of these dimensions represents the temporal aspect of the datasets, which we wish to preserve (because we wish to know what activity is taking place at any instant of time), the hidden feature we seek to uncover can be searched along one dimension only.

In correlating the inertial measurements with the ECG measurements, we face one challenge. The inertial sensors encode the robustness of the activity in terms of both the frequency and amplitude of the measurement they produce whereas the ECG encodes the same activity mainly along the frequency dimension (the heart beats faster when the body exerts). In order to cohesively correlate these aspects, we can use spectrum analysis. Fourier based spectrum analysis, however, reveals the spectral aspects of a signal but in doing so, completely hides its temporal aspects. For the analysis of ECG, this can be problematic. For example, a specialist may determine that the heart was exerting, may not, however, be able to determine why or when. The continuous wavelet transform [14] can be applied to capture both temporal and spectral aspects in the same dataset. Hence, for each of the channels we have, the wavelet transform results in a matrix whose rows and columns encode the frequency and temporal components, respectively.

Each of the channels we consider is unreliable when considered independently. Indeed, a closer look at Fig. 1 reveals that all the ECG channels suffer from a strong, long-term drift and zooming into a single ECG channel reveals a considerable distortion by motion artefacts. Putting the measurements of all the channels together along with their continuous wavelet transform enables to model a single activity as a three-way tensor (i.e., as a three dimensional array) as shown in Fig. 2.

B. Tensor Decomposition

In the same way a matrix can be decomposed (factorised) into basic constituting elements, a tensor can be decomposed into basic constituting elements. However, unlike decomposing a matrix, decomposing a tensor is not straightforward. To start with, an assumption has to be made as regards the number of the hidden factors, whereas this is done automatically with SVD. Secondly, a tensor has to be unfolded (or flattened) into a matrix before it can be decomposed, which influences the outcome of the decomposition.

A closer look into the activity tensor in Fig. 2 reveals that it provides three complementary views which can serve

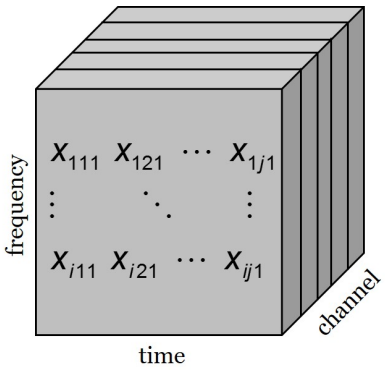

Fig. 2. A three-way tensor for modelling the activity of a person in terms of the measurements of a wireless ECG and inertial sensors along with their continuous wavelet transform.

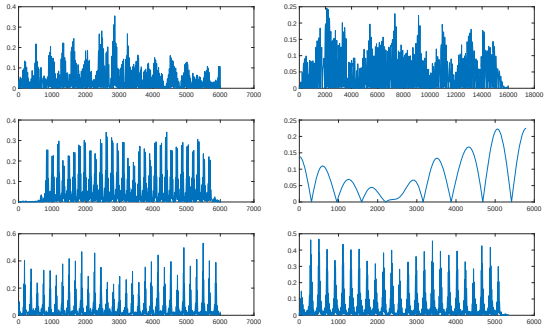

Fig. 3. One of the four factors of the **B** matrix (temporal aspects) of the tensor decomposition describing the activities of a person (Left to right and top to bottom: *cycling*, *climbing* up and down a staircase, *jumping* up and down on a flat surface, *push-ups*, *running*, and *skipping*).

different purposes. For example, the front view provides a matrix describing a single channel. This view is called the front slice. The top view provides a matrix describing the temporal aspect of all the channels for a single frequency. In other words, if a given activity manifested itself in terms of a single frequency component, it is easier to determine at which time this happened and whether all the sensors have registered this event. This view is called the horizontal slice. Likewise, the side view provides a matrix describing the manifestation of all frequency components in all the channels at any given time. This is called the lateral slice. It is this flexibility, among others, which makes a tensor decomposition desirable. The chief task of a tensor decomposition is to identify multidimensional features in terms of which activities can be categorized. Compared to the size of the tensor, the basic features are significantly small in size, so that the clustering process is computationally tractable.

A tensor analysis begins by unfolding (flattening) the tensor into a matrix. The unfolding can take place in different ways, but whichever way is chosen, the entries along each dimension form a column vector. Suppose, the first three front slices of the original tensor[2] are represented as follows:

$$\mathbf{X}_1 = \begin{bmatrix} 12 & 10 & 0 & \cdots \\ 8 & 2 & 120 & \cdots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix}$$
$$\mathbf{X}_2 = \begin{bmatrix} 15 & 10 & 0 & \cdots \\ 12 & 4 & 60 & \cdots \\ \vdots & \vdots & & \vdots \end{bmatrix} \quad (5)$$
$$\mathbf{X}_3 = \begin{bmatrix} 26 & 8 & 21 & \cdots \\ 50 & 3 & 15 & \cdots \\ \vdots & \vdots & \vdots & \cdots \end{bmatrix}$$

[2]This corresponds to the wavelet transform of the first three channels (say, the outputs of the X, Y, and Z axes of the accelerometer). The wavelet transform of a channel, as we already highlighted, is an $n \times m$ matrix.

The mode-1 unfolding of the above tensor takes each column vector of the matrices as they are and put them together side-by-side:

$$\mathbf{X}_{(1)} = \begin{bmatrix} 12 & 10 & 0 & \cdots & 15 & 10 & 0 & \cdots & 26 & \cdots \\ 8 & 2 & 120 & \cdots & 12 & 4 & 60 & \cdots & 50 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} \quad (6)$$

The mode-2 unfolding takes each raw entry of the matrices and places them as column vectors in a single matrix:

$$\mathbf{X}_{(1)} = \begin{bmatrix} 12 & 8 & \cdots & 15 & 12 & \cdots & 26 & 50 & \cdots \\ 10 & 2 & \cdots & 10 & 4 & \cdots & 8 & 3 & \cdots \\ 0 & 120 & \cdots & 0 & 60 & \cdots & 21 & 15 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} \quad (7)$$

Likewise, the mode-3 unfolding takes each lateral entry (along the time dimension) of the matrices and places them as column vectors in a single matrix. Afterwards, the unfolded matrix can be decomposed as if it were a normal matrix. There are different tensor decomposition strategies, but we employ the canonical polyadic decomposition (referred in the literature as CANDECOM/PARAFAC, or, in short, CP) [15] which decomposes a three-way tensor into three matrices:

$$\mathcal{X} = \mathbf{ABC} \quad (8)$$

or

$$\mathcal{X} = \sum_{r=1}^{R} \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r \quad (9)$$

where $\mathbf{a}_r$, $\mathbf{b}_r$, and $\mathbf{c}_r$, are the r-th columns of the matrices **A** **B**, and **C**, respectively. In the existence of a strong correlation, the activity tensor can be approximated only by the outer product of the first K column vectors of the matrices **A**, **B**, and **C**, respectively. The three basic matrices have the following significance: The matrix **A** relates the frequency components of the activity with the hidden features. The matrix **B** associates the unique features with the temporal aspect of the activity and the matrix **C** associates the different channels with the hidden features.

V. Evaluation

The focus of this section is on the scope and usefulness of a tensor decomposition in modelling and reasoning about different physical activities. We sample all the sensors at 512 samples per second. Then we apply the continuous wavelet transform on the row data in order to associate spectral aspects with temporal aspects. This results in a single channel being modelled by a matrix whose rows and columns express spectral and temporal aspects, respectively. The organisation of the matrices into a multi-dimensional array results in an activity tensor.

The CANDECOMP/PARAFAC (CP) requires as its input a tensor and the number of basis factors in terms of which the row data can be explained. In order to determine the number of basis factors, we began with two, performed the tensor decomposition, and evaluated the accuracy with which the reverse process produced the original tensor using Equation 5. For all the activities we considered in our analysis, no appreciable improvement in the *coefficient of determination* (R^2) was observed when we increased the basis factors beyond four. Therefore, we set the number of factors to four.

As we already mentioned in the previous section, the CP decomposition yields the **A**, **B**, and **C** matrices. Since our main concern is discriminating between different activities and not as such in determining which of the channels make the most contribution for the discrimination, the matrices **A** and **B** are the most important for our further analysis.

Fig. 3 displays how one of the four basis factors of the **B** matrix describes six different activities. Since the **B** matrix encodes the temporal aspect of an activity, the basis factor represents one unique temporal aspect. From visual inspection, it can be seen that the different activities have different footprints. If we take a small time window ΔT and add the square of the samples, we get the energy of the activity for that time window. Computing the energy this way by sliding the window and normalising the overall result by the number of windows, we get the average energy:

$$\frac{1}{N}\sum_{n\in N}\left(\sum_{k=1}^{T} a_k^2\right) \tag{10}$$

From this factor alone, it is possible to discriminate between the different activities. Nevertheless, not all the activities can be recognised with equal confidence. Whereas the *push-up* activity can be recognised with hundred percent accuracy, this cannot be achieved for discriminating between cycling and jumping as well as running and skipping. If we combine the first two factors of the **B** matrix, the recognition accuracy improves slightly, but not significantly. Other approaches can be adopted to systematically combine all the four temporal factors in order to improve the recognition accuracy, but this will introduce complexity to both the modelling and the computation process.

Unlike the factors expressing the temporal aspects, the factors expressing spectral aspects do not uniquely discriminate between activities when considered independently. When considered collectively, however, their expression power is much superior than the temporal aspects. First of all, the spectral maximum of a given basis factor can be determined uniquely for each activity. This maximum does not change appreciably from experiment to experiment. An exception to this is the *cycling* activity. Secondly, when the f_{max} of the different basis factors are added, they uniquely express an activity. This is despite the fact that some activities have similar maxima in one or more basis factors.

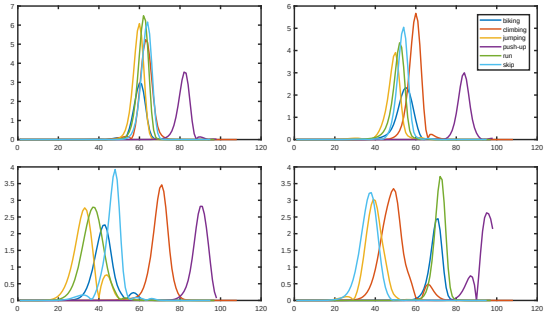

Fig. 4. The factors of the **A** matrix (spectral aspects of the tensor decomposition) describing the activities of a person in one of the experiments.

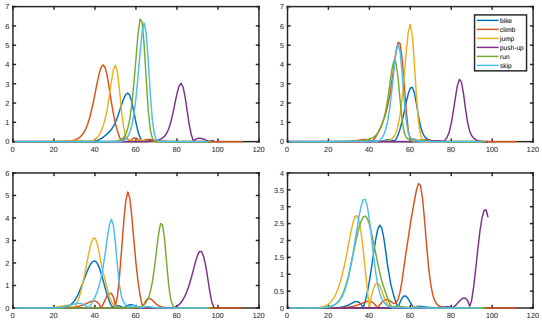

Fig. 5. The factors of the **A** matrix (spectral aspects of the tensor decomposition) describing the activities of a person in a different experiment.

In order to explain these two aspects, we refer to Figures 4 and 5 where we plotted the four spectral basis factors (the columns of the **A** matrix of the CP decomposition) for all the activities. The two figures correspond to two different experiments involving the same sets of activities. The colour of a plot corresponds to a particular activity. Summing the square of the magnitudes of all the factors for a particular activity amounts to obtaining the energy spectral density of the signal for that particular activity. We realise that the energy spectral density is not unique enough to discriminate between the activities, particularly when they involve intense movements. If, on the other hand, we obtain the frequencies at which the spectral density is maximum in the basis factors and simply add them, we can uniquely determine to which of the activities a particular sum refers.

In order to determine how the sum of frequencies differ from experiment to experiment for any given activity, we conduct six different experiments for each activity and examine the difference using Equation 11:

TABLE I
THE FREQUENCY SIGNATURE OF THE DIFFERENT MOVEMENTS FOLLOWING A TENSOR DECOMPOSITION.

Activity	Cyle	Climb	Jump	Push-up	Run	Skip
$\sum_{\substack{i,j=1 \\ i \neq j}}^{N} \lvert S_i - S_j \rvert$	35	17	1	15	3	2
$\frac{1}{N} \sum_{i=1}^{N} f_{max}^{i}$	226	228	183	352	225	203

TABLE II
THE FREQUENCY DIFFERENCE BETWEEN THE MAXIMUM FREQUENCIES OF THE DIFFERENT MOVEMENT TYPES FOLLOWING A TENSOR DECOMPOSITION.

	Cycle	Climb	Jump	Push-up	Run	Skip
Cycle	-	2	43	126	1	23
Climb	2	-	45	124	3	25
Jump	43	45	-	169	42	20
Push-up	126	124	169	-	127	149
Run	1	3	42	127	-	22
Skip	23	25	20	149	22	-

$$\Delta E = \sum_{\substack{i,j=1 \\ i \neq j}}^{N} |S_i - S_j| \tag{11}$$

where S_i and S_j refer to the sum of frequencies in experiment i and j, respectively. Tab. I lists the result. Ideally, the difference should be zero if the sum of frequencies uniquely expresses an activity. This, however, is not the case in reality. In order to appreciate the discrimination power of this approach, we consider the sums of frequencies of the different activities. The third row of Tab. I lists the average sum of frequencies for each activity (averaged over six different experiments). Tab. II displays the distance between the sums of frequencies for the different activities. Except for *cycling*, *climbing*, and *running*, discriminating between all the other activities results in almost $100\,\%$ accuracy. A false positive can be avoided for *running* almost with $100\,\%$ accuracy, since its sum of frequencies can be determined with great confidence, unlike for *climbing* and *cycling* whose sums of frequencies vary appreciably from experiment to experiment.

VI. CONCLUSION

In this paper we expressed the measurements of inertial sensors and a wireless electrocardiogram as elements of a three-way tensor in order to model different human activities. The inertial sensors consist of a 3D accelerometer and a 3D gyroscope and measure three dimensional rectilinear and curvilinear accelerations, respectively, whereas the wireless ECG provides three complementary channels (left leg-left arm, left leg-right arm, and left arm-right arm) measuring cardiac action potentials. Thus, we decomposed the activity tensor into three matrices and analysed the expression power of the factors encoded in these matrices. The first matrix encodes spectral aspects whilst the second matrix encodes temporal aspects. The third matrix was not of interest. We demonstrated how easily the basis factors describing spectral aspects can be used to recognise activities. One way to improve the expression power of the spectral aspects of a tensor decomposition will be to use a Bayesian Network, so as to associate a probabilistic value to each of the spectral aspects. This will be the focus of our future work.

REFERENCES

[1] W. Dargie, *Principles and Applications of Ubiquitous Sensing*. John Wiley & Sons, 2017.

[2] C. J. Deepu, X. Zhang, C. H. Heng, and Y. Lian, "A 3-lead ecg-on-chip with qrs detection and lossless compression for wireless sensors," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 63, no. 12, pp. 1151–1155, 2016.

[3] Z. C. Haberman, R. T. Jahn, R. Bose, H. Tun, J. S. Shinbane, R. N. Doshi, P. M. Chang, and L. A. Saxon, "Wireless smartphone ecg enables large-scale screening in diverse populations," *Journal of cardiovascular electrophysiology*, vol. 26, no. 5, pp. 520–526, 2015.

[4] A. Burns, B. R. Greene, M. J. McGrath, T. J. O'Shea, B. Kuris, S. M. Ayer, F. Stroiescu, and V. Cionca, "Shimmer–a wireless sensor platform for noninvasive biomedical research," *IEEE Sensors Journal*, vol. 10, no. 9, pp. 1527–1534, 2010.

[5] W. Dargie, "A stochastic fusion technique for removing motion artefacts from the measurements of a wireless ecg," in *Multisensor Fusion and Integration for Intelligent Systems (MFI), 2016 IEEE International Conference on*, pp. 449–454, IEEE, 2016.

[6] A. Fratini, M. Sansone, P. Bifulco, and M. Cesarelli, "Individual identification via electrocardiogram analysis," *Biomedical engineering online*, vol. 14, no. 1, p. 78, 2015.

[7] M. Li, V. Rozgić, G. Thatte, S. Lee, A. Emken, M. Annavaram, U. Mitra, D. Spruijt-Metz, and S. Narayanan, "Multimodal physical activity recognition by fusing temporal and cepstral information," *IEEE transactions on neural systems and rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 18, no. 4, p. 369, 2010.

[8] J.-B. Martens, "The hermite transform-theory," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, no. 9, pp. 1595–1606, 1990.

[9] S. Rahman, "Wiener–hermite polynomial expansion for multivariate gaussian probability measures," *Journal of Mathematical Analysis and Applications*, vol. 454, no. 1, pp. 303–334, 2017.

[10] F. Massé, M. V. Bussel, A. Serteyn, J. Arends, and J. Penders, "Miniaturized wireless ecg monitor for real-time detection of epileptic seizures," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 12, no. 4, p. 102, 2013.

[11] F. Rincón, P. R. Grassi, N. Khaled, D. Atienza, and D. Sciuto, "Automated real-time atrial fibrillation detection on a wearable wireless sensor platform," in *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pp. 2472–2475, IEEE, 2012.

[12] C. Roselli, M. D. Chaffin, L.-C. Weng, S. Aeschbacher, G. Ahlberg, C. M. Albert, P. Almgren, A. Alonso, C. D. Anderson, K. G. Aragam, *et al.*, "Multi-ethnic genome-wide association study for atrial fibrillation," *Nature genetics*, p. 1, 2018.

[13] "Mit-bih arrhythmia database," 2001.

[14] G. G. Walter and X. Shen, *Wavelets and other orthogonal systems*. CRC press, 2018.

[15] A. Stegeman and N. D. Sidiropoulos, "On kruskals uniqueness condition for the candecomp/parafac decomposition," *Linear Algebra and its applications*, vol. 420, no. 2-3, pp. 540–552, 2007.