Citation/Reference	Billiet L, Swinnen T, Westhovens R, de Vlam K, Van Huffel S **Activity Recognition for Physical Therapy – Fusing Signal Processing Patterns and Movement Patterns** 3rd international Workshop on Sensor-based Activity Recognition and Interaction, 2016
Archived version	Author manuscript: the content is identical to the content of the published paper, but without the final typesetting by the publisher
Published version	http://dx.doi.org/10.1145/2948963.2948968
Journal homepage	http://iwoar.org/2016/
Author contact	Lieven.billiet@esat.kuleuven.be *+ 32 (0)16 327685*
IR	

(article begins on next page)

Activity Recognition for Physical Therapy

Fusing Signal Processing Features and Movement Patterns

[1]**Lieven Billiet**[*†]
KU Leuven
ESAT - STADIUS
Kasteelpark Arenberg 10, 2446
3001 Leuven, Belgium
lieven.billiet@kuleuven.be

[1]**Thijs Swinnen**[‡§¶]
University Hospitals Leuven
Division of Rheumatology
Herestraat 49
3000 Leuven, Belgium
thijs.swinnen@uzleuven.be

Rene Westhovens[‡§]
University Hospitals Leuven
Division of Rheumatology

Kurt de Vlam[‡§]
University Hospitals Leuven
Division of Rheumatology

Sabine Van Huffel[*†]
KU Leuven
ESAT - STADIUS

ABSTRACT

This paper discusses the classification of activities in the context of physical therapy. Usually, specific standardized activities are subjectively assessed, often by means of a patient-reported questionnaire, to estimate a patient's activity capacity, defined as the ability to execute a task. Automatic recognition of these activities is of vital importance for a more objective and quantitative approach to the problem. The proposed accelerometry-based algorithm merges standard signal processing features with information obtained from direct activity pattern matching using dynamic time warping (DTW) in a linear model. This study with 28 spondyloarthritis patients performing 10 activities shows the improvement in activity classification accuracy due to the fusion of the two approaches, up to 93.6%. This is a significant increase compared to similar models based on either of the approaches alone ($p < 0.01$). Although this paper mainly contributes to the activity recognition step, it also briefly discusses the advantages of the approach with regard to further automated evaluation of the recognized activities.

Author Keywords

Human activity recognition; quantified rehabilitation; pattern matching; dynamic time warping; time and frequency features; early fusion; LDA.

[1]LB and TS contributed equally.

[*]KU Leuven, Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics.

[†]iMinds Medical IT, Leuven

[‡]University Hospitals Leuven, Division of Rheumatology.

[§]KU Leuven, Department of Development and Regeneration, Skeletal Biology and Engineering Research Center.

[¶]KU Leuven, Department of Rehabilitation Sciences, Musculoskeletal Rehabilitation Research Unit

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

iWOAR '16, June 23-24, 2016, Rostock, Germany.

© 2016 ACM. ISBN 978-1-4503-4245-2/16/06...$15.00

DOI: http://dx.doi.org/10.1145/2948963.2948968

INTRODUCTION

In rehabilitation, a physical therapist is often interested in the activity capacity of a patient i.e. the ability to execute a task. One way of quantifying this capacity is the use of subjective scoring systems. A typical example is the Bath Ankylosing Spondylitis Functional Index (BASFI), geared towards axial spondyloarthritis patients [5]. It is a questionnaire listing a number of typically limiting activities (sitting, standing, reaching up etc.) in this group of patients. The patient indicates the ability to perform each activity using a score on a numerical rating scale. During treatment, one expects to observe a positive evolution of the capacity. However, such self-reported scoring does not always reflect the clinical impression of the treating physical therapist or rheumatologist and has been shown to be influenced by psychosocial factors. The resulting variability may lead to invalid appraisals of activity capacity and obscure the patient's progress during treatment [4]. As an objective alternative, the activities could be measured and quantified with clearly defined technology-based capacity features. For example, one could use 3D motion capture to obtain kinematic information [8]. Yet, such a system requires a motion lab for precise measurements or, at least, a stereoscopic camera. As an alternative, one could use one or a few wearable sensors (e.g. accelerometers) to track the motion pattern. In this case, the aforementioned activities can even be done in the home environment, allowing for more regular follow-up data. Such an approach requires automatic recognition of the motions and (activity-specific) parameter extraction. The focus of this paper is on the first aspect, the recognition of the activities, rather than their evaluation.

In the last decades, accelerometers have been widely applied for activity recognition [2, 3, 16]. The algorithms can be divided in two categories: window-based or template-based. Most of the approaches are in the first category, although some overlap can be observed. Sliding windows, either of fixed or variable size, create a discrete problem since every window can be assigned a class label. It is particularly suited for long term study of repetitive or static activities. These can be characterized by features extracted from each window in the time domain, frequency domain or time-frequency (e.g. wavelets). In the remainder, this kind of features will be called Signal Processing (SP) features. They have been used

successfully, e.g. a study with 7 healthy subjects performing 6 basic activities classified 6s activity windows with high accuracy [9]. Yet, it has been shown that patient populations exhibit a greater variability than healthy control subjects e.g. for gait patterns [11]. As an example, a smartphone activity recognition system for 5 activities in a cross-subject setting observed a performance drop of more than 10% (to about 75%) when comparing healthy control subjects with Parkinson patients [1].

The template-based approach has attracted less attention. At the core, it uses class templates (e.g. accelerometer data segments). Mostly, the activities under study are transitory rather than static. The templates are compared to new data to infer the unknown class labels. The most obvious drawback of the method is the inability of static patterns to capture variability. As a result, the approach is sometimes restricted to patterns with limited variability e.g. heartbeat classification [15]. Another option is not matching the entire activity patterns, but focus on subpatterns. For example, Zhang and Sawchuk define motion primitives for a bag-of-features approach to classify 9 activities [18]. If one matches the whole pattern, the set of activities is often limited and a flexible matching such as dynamic time warping (DTW) is employed to allow for some variability. This further allows for some robustness with regard to differences in sensor placement [12]. In a study to detect sit-to-stand and stand-to-sit transitions using DTW, Ganea *et al*. obtain accuracies up to 95% (healthy subjects) and 89% (chronic pain patients) [7]. Another activity-limited study focuses on derivative DTW (DDTW) to accurately classify three gait-related activities.

Recent work by Margarito *et al*. compares statistical learning classifiers to several template matching techniques and decides only slightly in favor of the former. Pattern matching accuracy is around 75-80% for 8 main classes, whereas the optimal neural network could attain on average 85% accuracy [13].

Lastly, some commercial products should also be mentioned. For example, the MoveTest (McRoberts [14]) and the Valedo (Hocoma AG [10]) have been designed to support physiotherapists and patients. The MoveTest focuses only on a small subset of relevant tasks (e.g. sit-to-stand and timed up-and-go). The Valedo is geared towards low back pain and provides an interface to 45 exercises at home via serious gaming. However, it seems to focus mostly on changes of orientation, reaction speed and range of motion in the context of the specific gamified exercise with tightly controlled movements. Yet, the whole patient-specific movement trajectory could be of interest.

To summarize, the current state-of-the-art mostly uses sliding window approaches for activity recognition and demonstrates good accuracy on healthy subjects. Less frequent pattern matching approaches are less powerful and made viable by restricting the type and number of activities, among other constraints. In contrast, this study focuses on the evaluation of patients performing a variety of clinically important activities during their routine follow-up by a physical therapist. It explicitly combines the concepts of sliding window features and activity patterns to obtain superior performance. The developed algorithm also aims at facilitating capacity feature extraction, although this is not the main focus.

Abbreviation	Description
getup	getting up starting from lying down
liedown	lying down starting from stance
maxreach	reaching up as far as possible
pen	picking up a pen from the ground
pen5	repeating pen 5 times, as fast as possible
reach5	touching a mark 5 times, as fast as possible
STS	performing a sit-to-stand movement
STS5	repeating STS 5 times, as fast as possible
sock	sitting down and putting on socks
stairs	climbing a stairs (13 steps)

Table 1. Description of the activities

The remainder of the paper is structured as follows. The next Section discusses the patient study setup and outlines the classification algorithm that has been developed. Classification results are presented in the subsequent results Section. It is followed by a discussion Section. Finally, we present our conclusions and perspectives for future work.

DATA AND METHODS

A patient study was performed at the Division of Rheumatology of the University Hospitals, Leuven [17]. The study protocol was approved by the Medical Ethics Committee (ML 5236). It will be discussed first. Afterwards, the developed algorithm for activity classification will be presented in detail, followed by a description of the classification setup.

Test protocol

Data collection involved 28 subjects (16 male, 12 female) diagnosed with axial spondyloarthritis (aSpA). This rheumatic disease is characterized by inflammation of the spinal region. It often induces spinal fusion (ankylosis). The selected patients' functional capacity, quantified by the BASFI score, ranges from 0/10 (best) to 8.1/10 (worst), with an average score of 3.14/10. They are on average 43.69 years old (standard deviation 10.45).

The experiments took place in a controlled laboratory environment. The patients were equipped with a two-axial accelerometer (Sensewear Pro 3 Armband, Bodymedia Inc, Pittsburgh, USA) sampling at 32Hz. It was mounted on the biceps of the dominant arm, its orientation matching the longitudinal and transversal axes. Only one sensor was used to limit the inconvenience for the patient. The upper arm is a good location to measure the movement of the entire body, since it is still relatively close to the center of gravity, while at the same time yielding information about arm movements. Next, the patients performed a series of 10 activities inspired by the BASFI tasks. All of them are transitions, since these reveal most information on activity capacity. They are listed in Table 1. Some of them are fast repetitions of others, partly because this leads to more (subject-specific) standardized behavior and partly to study the effect of speed. Nevertheless, they will be treated and characterized as separate activities. Every subject performed each activity twice. Hence, in total, 560 activity trials were acquired, 56 for each class.

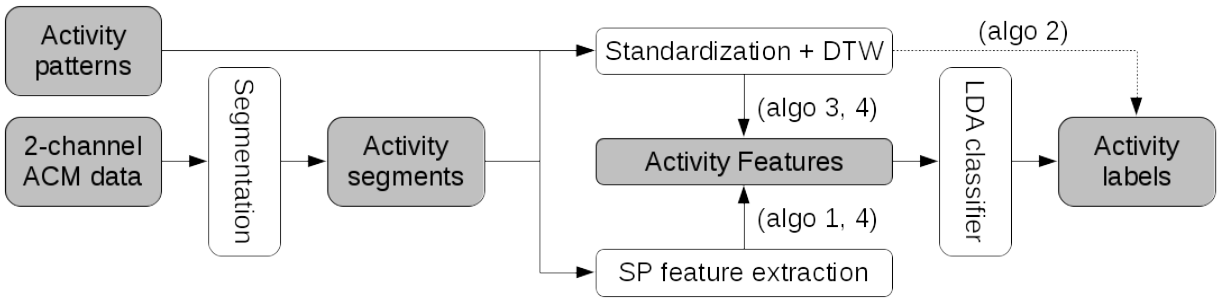


Figure 1. Flow chart of the various algorithms. (1) uses only SP features with LDA, (2) decides directly based on DTW, (3) uses only DTW features with LDA and (4) fuses SP and DTW features for use with LDA

Classification approach

This study consists of four classification approaches (referred to as 'algorithms'), presented in the flow diagram in Figure 1. Grey blocks indicate data, white blocks represent algorithm building blocks. We can distinguish four of them: Segmentation, SP feature extraction, DTW and the Linear Discriminant Analysis (LDA) classifier.

Segmentation

Activities are recorded continuously. In between the actual activities, a patient simply remains standing. Hence, the accelerometer signals are static, except for the occurrence of activities. Based on these assumptions Milosevic M. developed a multi-pass dynamic threshold method to extract the activities (semi-)automatically (see [17]). This study adopts this algorithm to obtain activity segments, additionally filtering the signal with a butterworth filter (4th order, cutoff at 3.2Hz).

As Figure 1 shows, all developed algorithms pass through the segmentation stage. 2-channel accelerometry data (ACM) is transformed into (2-channel) segments for further processing. Overall, this yields a 2-stage segmentation-recognition approach. It is only viable in a controlled setup, due to the limited movement in between activities. Yet, it is possible to have the patient perform the same procedure in the home environment.

SP feature extraction

Once segments have been obtained, the standard approach used in window-based algorithms can be applied. Only, instead of defining features for a sliding window they can be extracted directly from the activity segments, since these are the only regions of interest. The following features have been selected (per channel), many of them common in literature:

- **mean** (2 features): mostly important to distinguish poses or pose changes.
- **standard deviation** (2 features): captures the spread of the data.
- **3 subpart means** (6 features): three uniform bin means are a coarse representation of signal shape. Due to its coarseness, it is somewhat robust against the trial-to-trial and subject-to-subject variability.
- **power** (2 features): a measure for the intensity of the activity.
- **power spectral entropy (PSE)** (2 features): yields information on the spectral structure or complexity of the signal.
- **autocorrelation** (2 features): allows to separate repetitive and non-repetitive tasks as it is an indicator of self-similarity.

Hence, in total, 16 features are used to characterize a trial.

LDA classifier

Once the features have been calculated, one can train a classifier and predict the class of newly acquired samples. Since the amount of data is relatively small for 10 classes, a linear classifier was selected. It has the advantage of model simplicity, which limits the risk of overfitting, particularly important with a limited data set. Linear discriminant analysis fits a multivariate normal distribution to each class. In the linear case, the classes all share the same covariance matrix. Even though these conditions (multivariate normality, equal covariance) are often not met, the resulting discriminant class boundaries yield good results in practice [6]. The method has the additional advantage of a natural extension to multiple classes.

DTW

Dynamic time warping matches signals by a nonlinear deformation (warping). In essence, it matches points of the two shapes, as indicated in Figure 2. Based on these correspondences, the signals can be aligned. In this study, the objective function of the warping is the minimization of the euclidian distance between the matched and deformed signals. Because our primary concern is shape rather than amplitude

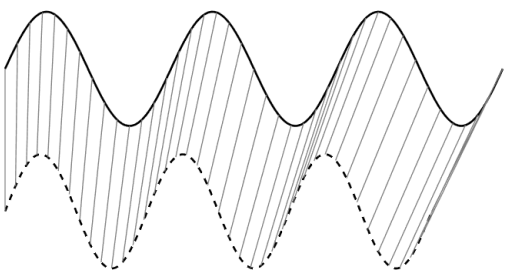


Figure 2. Example of DTW matching of two sinusoids of different frequency.

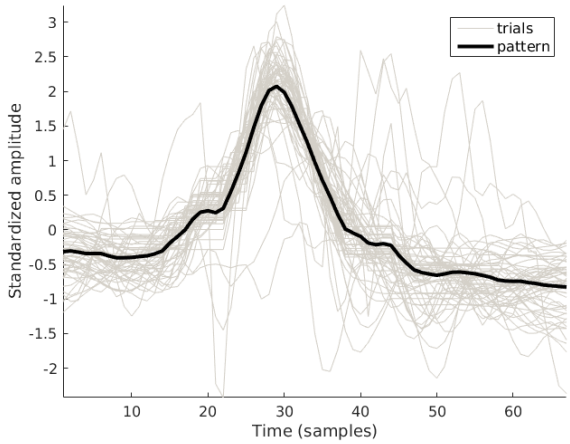


Figure 3. Pattern extraction for STS, longitudinal axis

differences, the signals are normalized (standardized) first. The DTW matching is calculated using the implementation of Zhou and la Torre [19]. It is an extension of the standard DTW: it includes multidimensional matching (multiple channels) and multiple alignment (more than two signals).
In our approach, activity patterns are first generated using DTW on the standardized training data, one for each class. New data can be compared by matching it to the stored patterns. In the simplest approach (algorithm 2 in Figure 1) the class pattern with the lowest DTW matching objective value (euclidian distance) indicates the predicted class. In a more elaborate approach (algorithm 3 and 4 in Figure 1) the objective values for DTW matchings to all stored class patterns are used as features in the LDA classifier. This assumes similarity to a class pattern might reveal information, even if the new sample does not belong to that class.

Algorithms

The flow in Figure 1 shows four algorithms. All start with segmentation.

- **algo 1 (linSP)** At training time, SP features are extracted from training segments and used to train an LDA classifier. This classifier can be used for prediction.
- **algo 2 (purePat)** DTW aligning of the standardized training segments yields activity patterns. Figure 3 gives an example for STS. New (standardized) segments can be matched against these patterns. The most similar pattern, expressed by the euclidian distance, indicates the class.
- **algo 3 (linPat)** Algo 3 uses the same activity patterns as algo 2. Next, each training segments is matched against all patterns. This creates a new training set: the resulting euclidian distances to each pattern are the training features. The LDA classifier is trained with this new training set. Test segments are converted to the feature space by DTW to the activity patterns and fed into the LDA classifier for prediction.
- **algo 4 (linPatSP)** Algo 4 merges algo 1 and algo 3 by means of early fusion. The feature data (SP features and pattern distances) are concatenated. The LDA classifier is trained on this extended training data set.

Classification setup details

The evaluation of the algorithms is structured in a leave-one-subject-out (LoSo) sense. Hence, all trained classifiers are subject-independent. For each evaluation, 504 training samples and 56 test samples (corresponding to one subject) are used. The performances of the four algorithms are compared via a paired t-test.

RESULTS

The LoSo accuracy of all algorithms for all subjects is displayed in Table 2 (chance level is 0.1). The last rows show the average accuracy and the standard deviation. The best algorithm for each subject is highlighted in bold. Based on these results, some conclusions can be drawn, backed up by statistical analysis. The pure pattern matching clearly performs worst, though not in all cases. It has a moderate average accuracy with a high standard deviation. Its accuracy is lower than linSP ($p = 0.0194$) and significantly lower than linPat ($p = 6e\text{-}4$). The latter indicates similarity to multiple patterns is more conclusive than only looking for the closest match. Comparison of linSP and linPat is not conclusive ($p = 0.725$). Indeed, both their average accuracies and standard deviations are similar as well. Finally, linPatSP significantly

	linSP	**purePat**	**linPat**	**linPatSP**
subj1	0.65	0.85	0.85	**1**
subj2	0.9	0.7	0.8	**0.95**
subj3	0.7	**0.95**	**0.95**	0.9
subj4	**0.8**	0.6	0.7	**0.8**
subj5	0.8	0.95	**1**	**1**
subj6	**0.9**	0.4	0.7	0.85
subj7	**0.85**	0.55	0.8	0.8
subj8	0.9	0.75	0.85	**1**
subj9	0.8	0.9	**0.95**	0.9
subj10	**0.9**	0.8	0.8	**0.9**
subj11	**0.95**	0.3	0.5	0.85
subj12	0.95	0.85	**1**	**1**
subj13	**1**	0.9	0.8	**1**
subj14	**1**	0.9	0.95	**1**
subj15	**1**	0.95	**1**	**1**
subj16	0.85	**1**	0.95	**1**
subj17	0.75	0.55	0.75	**0.8**
subj18	**0.95**	0.8	0.85	**0.95**
subj19	0.9	0.95	**1**	**1**
subj20	0.9	0.9	0.95	**1**
subj21	0.9	0.45	0.8	**0.95**
subj22	0.9	0.85	0.95	**1**
subj23	0.9	0.8	0.9	**1**
subj24	**1**	0.45	0.75	0.9
subj25	0.95	0.85	0.95	**1**
subj26	0.4	0.55	0.55	**0.7**
subj27	**1**	0.85	0.9	**1**
subj28	0.75	0.9	**1**	0.95
avg	0.866	0.759	0.856	**0.936**
std	0.131	0.196	0.133	**0.084**

Table 2. LoSo accuracies of all algorithms for all subjects.

	getup	liedown	maxreach	pen	pen5	reach5	STS	STS5	socks	stairs
getup*	54	0	0	0	0	0	0	0	0	0
liedown*	0	56	1	0	0	0	0	0	0	0
maxreach*	0	0	55	0	0	0	0	0	0	0
pen*	0	0	0	48	0	0	4	0	1	0
pen5*	0	0	0	0	51	0	0	0	0	0
reach5*	0	0	0	0	2	56	0	0	0	0
STS*	0	0	0	6	0	0	52	0	1	0
STS5*	0	0	0	0	0	0	0	45	0	2
socks*	2	0	0	2	1	0	0	0	54	1
stairs*	0	0	0	0	2	0	0	11	0	53

Table 3. Accumulated confusion matrix over all subjects for linPatSP. Rows are estimated* classes, columns actual classes.

outperforms all others (p values of 0.0021, 1.4e-6 and 5.2e-5, respectively). Although its performance is comparable to linSP for many subjects, its lower standard deviation shows it is able to be more robust to variability by fusion of complementary information. This is particularly clear for subject 26: though all methods have low performance, the fusion allows to obtain an acceptable accuracy of 70%.

Table 3 shows the pooled confusion matrix. The class-averaged sensitivity is 93.6% (std 6.4%). The specificity is 99.3% (std 0.8%). As can be observed, liedown and reach5 are always identified correctly (perfect sensitivity). STS5 has the lowest sensitivity, 80.4%, because it is often mistakenly identified as climbing the stairs. Recalling the sensor is mounted on the arm, the movements can seem similar. The templates for the movements are similar as well, but different enough to be distinguished by a human observer. However, they generate similar statistics. Experiments show that indeed, linSP performs worse than linPat or linPatSP for the STS activity. The confusion matrix further shows the decision boundary favors stairs over STS5, since stairs are not that often misclassified as STS5. The same argument can be applied for the confusion between pen and STS: the shapes are more dissimilar than the SP features. Yet, for both cases, combining SP features and patterns (linPatSP) still yields an advantage (up to 5%) over either of the single classifiers.

DISCUSSION

All approaches in this work are subject-independent. Partly, this is a pragmatic choice. Since only two samples of each activity are available for any subject, even including half of it would only mean one data point (out of 56) would be the subject-dependent part. Yet, additional experiments show this already leads to a minor improvement, especially for subject 26. Another way to estimate the influence of between-subject variability is to check whether misclassifications are tied to subjects, that is, if both instances of the subjects are misclassified, rather than only one. In the latter case, the within-subject variability would be more important. As it turns out, this test is inconclusive: in 11 cases, both trials are misclassified; in 14 cases, only one of the trials of a subject is misclassified. On the other hand, one can also observe that in particular cases, bad performance of a subject is indeed tied to a misclassification of both trials. Generally speaking though, this claim cannot be proven. This is a positive observation: it entails the algorithm does not need specific patient training, but performs well for any patient.

The flow outline in Figure 1 reveals a bottleneck. Before recognition is possible, segmentation should have been performed. The current segmentation algorithm is only applicable in a controlled scenario. It depends on the absence of motion when no physical exercise is being performed. Depending on the scenario, this constraint might be a problem. Currently, the longer-term aim is to have patients perform exercises at home. For example, they could take some time every day to perform the protocol. In such a setting, the activities can be recognized and important informative parameters can be extracted. When the data is transferred to the therapist, he or she can follow the evolution of the parameters.
However, one can argue that patients, even at home, do not perform activities as they would do in uncontrolled conditions. Therefore, the aim might be to identify the activities in daily life. In such cases, another segmentation approach would be necessary.

As indicated in the introduction, explicit pattern matching has some disadvantages. Moreover, the abundant positive results in literature seem to suggest window-based approaches are well suited to solve the activity recognition problem. Then, why still use patterns? It is mostly a principled choice, rooted in the observed clinical reality. Physiotherapists reason on how a motion is being performed. When presented with accelerometer data, this is the way they look at it. Ideally, motion, rather than acceleration patterns, would be the primary way of presenting the information. But even with acceleration patterns, points of interest can be defined. E.g. how long does it take a patient to get up, what is the average acceleration when bending over etc. Clinicians can reason on patterns, particularly when these are complemented with video recordings. Therefore, the fusion approach developed in this work tries to combine the best of both worlds: patterns are being matched to retain the interpretability and identify the specific interesting parts of the data, but SP features leverage their classification performance.

CONCLUSIONS AND FUTURE WORK

In this work, we introduced a fusion approach to classify activities in the context of physical therapy. After introducing the problem and related work, we discussed the patient study and the building blocks of the algorithm: Segmentation, Signal Processing features, Dynamic Time Warping and Linear

Discriminant Analysis. The results section showed that the subject-independent classifier with early fusion of pattern and SP features outperformed classifiers from a single modality. Hence, the work is a step towards objectifying physical therapy based on acceleration patterns.

In the future, the segmentation should be further improved to also allow for recognition in daily situations. A possible approach is a coarse subdivision of the signal based on SP features, followed by a refinement based on the known motion patterns. As a second goal, the matching obtained using our approach should be processed for the next stage: activity capacity assessment. To this end, activity-specific features should be defined on the patterns. When matching the patterns to data, the feature values for the corresponding parts of the predefined pattern parts can be calculated. Finally, these obtained measures will be integrated in an interpretable assessment system for comparison to current assessment methods such as BASFI. We are currently developing a data-driven medical scoring system to that end.

Finally, data should be captured at the patients' home to assess the impact of the standardized clinical environment. Recently, such data has been gathered already.

ACKNOWLEDGEMENTS

This research was supported by: Bijzonder Onderzoeksfonds KU Leuven (BOF), Center of Excellence (CoE): PFV/10/002 (OPTEC); KULeuven IDO funding: #3E140722 Sensor-based Platform for the Accurate and Remote monitoring of Kine(ma)tics Linked to E-health (SPARKLE); Belgian Federal Science Policy Office: IUAP #P7/19/ (DYSCO, 'Dynamical systems, control and optimization', 2012-2017).

REFERENCES

1. Albert, M. V., Toledo, S., Shapiro, M., and Koerding, K. Using mobile phones for activity recognition in Parkinsons patients. *Front Neurol 3*, 158 (2012).
2. Attal, F., Mohammed, S., Dedabrishvili, M., Chamroukhi, F., Oukhellou, L., and Amirat, Y. Physical Human Activity Recognition Using Wearable Sensors. *Sensors 15*, 12 (2015), 31314–31338.
3. Avci, A., Bosch, S., Marin-Perianu, M., Marin-Perianu, R., and Havinga, P. Activity Recognition Using Inertial Sensing for Healthcare, Wellbeing and Sports Applications: A Survey. In *Proc ARCS* (2010), 1–10.
4. Brionez, T. F., Assassi, S., Reveille, J. D., Learch, T. J., Diekman, L., Ward, M. M., Davis, J. C., Weisman, M. H., and Nicassio, P. Psychological correlates of self-reported functional limitation in patients with ankylosing spondylitis. *Arthritis Res Ther 11*, 6 (2009), 1–9.
5. Calin, A., Garrett, S., Whitelock, H., Kennedy, L. G., O'Hea, J., Mallorie, P., and Jenkinson, T. A new approach to defining functional ability in ankylosing spondylitis: the development of the bath ankylosing spondylitis functional index. *J Rheumatol 21*, 12 (1994), 2281–2285.
6. Duda, R., Hart, P., and Stork, D. *Pattern Classification*. Wiley, 2001.
7. Ganea, R., Paraschiv-lonescu, A., and Aminian, K. Detection and Classification of Postural Transitions in Real-World Conditions. *IEEE Trans Neural Syst Rehabil Eng 20*, 5 (2012), 688–696.
8. Garrido-Castro, J. L., Medina-Carnicer, R., Schiottis, R., Galisteo, A. M., Collantes-Estevez, E., and Gonzalez-Navas, C. Assessment of spinal mobility in ankylosing spondylitis using a video-based motion capture system. *Manual Ther 17*, 5 (2012), 422 – 426.
9. Gupta, P., and Dallas, T. Feature Selection and Activity Recognition System Using a Single Triaxial Accelerometer. *IEEE Trans Biomed Eng 61*, 6 (2014), 1780–1786.
10. Hocoma AG. Back pain therapy with the Valedo THERAPY Concept. `https://www.hocoma.com/world/en/products/valedo-therapy-concept/`, accessed April 6, 2016.
11. Huisinga, J. M., Mancini, M., St. George, R. J., and Horak, F. B. Accelerometry Reveals Differences in Gait Variability Between Patients with Multiple Sclerosis and Healthy Controls. *Ann Biomed Eng 41*, 8 (2012), 1670–1679.
12. Kale, N., Lee, J., Lotfian, R., and Jafari, R. Impact of Sensor Misplacement on Dynamic Time Warping Based Human Activity Recognition Using Wearable Computers. In *Proc WH '12* (2012), 7:1–7:8.
13. Margarito, J., Helaoui, R., Bianchi, A. M., Sartor, F., and Bonomi, A. G. User-Independent Recognition of Sports Activities From a Single Wrist-Worn Accelerometer: A Template-Matching-Based Approach. *IEEE Trans Biomed Eng 63*, 4 (2016), 788–796.
14. McRoberts. MoveTest: Unobtrusive assessment of Physical Function under supervised conditions. `http://www.mcroberts.nl/products/mt/`, accessed April 6, 2016.
15. Nair, T. R. G., Geetha, A. P., and Asharani, M. Methodology For Detection of QRS Pattern Using Secondary Wavelets. *CoRR abs/1408.0452* (2014).
16. Preece, S., Goulermas, J., Kenney, L., Howard, D., Meijer, K., and Crompton, R. Activity identification using body-mounted sensors – a review of classification techniques. *Physiol Meas 30* (2009), R1–R33.
17. Swinnen, T., Milosevic, M., Van Huffel, S., Dankaerts, W., Westhovens, R., and De Vlam, K. Proof of concept : the novel instrumented Bath Ankylosing Spondylitis Functional Index (iBASFI) shows promising reliability and validity in the assessment of activity limitations in axial spondyloarthritis. *J Rheumatol (accepted)* (2016).
18. Zhang, M., and Sawchuk, A. A. Motion Primitive-based Human Activity Recognition Using a Bag-of-features Approach. In *Proc SIGHIT IHI '12* (2012), 631–640.
19. Zhou, F., and la Torre, F. D. Generalized time warping for multi-modal alignment of human motion. In *Proc IEEE CVPR12* (2012), 1282–1289.