SELFBACK—Activity Recognition for Self-management of Low Back Pain

Sadiq Sani, Nirmalie Wiratunga, Stewart Massie and Kay Cooper

Abstract Low back pain (LBP) is the most significant contributor to years lived with disability in Europe and results in significant financial cost to European economies. Guidelines for the management of LBP have self-management at their cornerstone, where patients are advised against bed rest, and to remain active. In this paper, we introduce SELFBACK, a decision support system used by the patients themselves to improve and reinforce self-management of LBP. SELFBACK uses activity recognition from wearable sensors in order to automatically determine the type and level of activity of a user. This is used by the system to automatically determine how well users adhere to prescribed physical activity guidelines. Important parameters of an activity recognition system include windowing, feature extraction and classification. The choices of these parameters for the SELFBACK system are supported by empirical comparative analyses which are presented in this paper. In addition, two approaches are presented for detecting step counts for ambulation activities (e.g. walking and running) which help to determine activity intensity. Evaluation shows the SELFBACK system is able to distinguish between five common daily activities with 0.9 macro-averaged F1 and detect step counts with 6.4 and 5.6 root mean squared error for walking and running respectively.

Keywords Intelligent Decision Support Systems · Medical Computing and Health Informatics · Machine Learning

S. Sani (✉) · N. Wiratunga · S. Massie · K. Cooper
Robert Gordon University, Aberdeen, UK
e-mail: s.sani@rgu.ac.uk

N. Wiratunga
e-mail: n.wiratunga@rgu.ac.uk

S. Massie
e-mail: s.massie@rgu.ac.uk

K. Cooper
e-mail: k.cooper@rgu.ac.uk

1 Introduction

Low back pain (LBP) is a common, costly and disabling condition that affects all age groups. It is estimated that up to 90 % of the population will have LBP at some point in their lives, and the recent global burden of disease study demonstrated that LBP is the most significant contributor to years lived with disability in Europe [5]. Non-specific LBP (i.e. LBP not attributable to serious pathology) is the fourth most common condition seen in primary care and the most common musculoskeletal condition seen by General Practitioners [11], resulting in substantial cost implications to economies. Direct costs have been estimated in one study as 1.65–3.22 % of all health expenditure [12], and in another as 0.4–1.2 % of GDP in the European Union [7]. Indirect costs, which are largely due to work absence, have been estimated as $50 billion in the USA and $11 billion in the UK [7]. Recent published guidelines for the management of non-specific LBP [3] have self-management at their cornerstone, with patients being advised against bed rest, and advised to remain active, remain at work where possible, and to perform stretching and strengthening exercises. Some guidelines also include advice regarding avoiding long periods of inactivity.[1]

SELFBACK is a monitoring system designed to assist the patient in deciding and reinforcing the appropriate physical activities to manage LBP after consulting a health care professional in primary care. Sensor data is continuously read from a wearable device worn by the user, and the user's activities are recognised in real time. An overview of the activity recognition components of the SELFBACK system is shown in Fig. 1. Guidelines for LBP recommend that patients should not be sedentary for long periods of time. Accordingly, if the SELFBACK system detects continuous periods of sedentary behaviour, a notification is given to alert the user. At the end of the day, a daily activity profile is also generated which summarises all activities done by the user over the course of the day. The information in this daily profile also includes the durations of activities and, for ambulation activities (such as moving from one place to another e.g. walking and running), the counts of steps taken. The system then compares this activity profile to the recommended guidelines for daily activity and produces feedback to inform the user how well they have adhered to these guidelines.

The first contribution of this paper is the description of an efficient, yet effective feature representation approach based on Discrete Cosine Transforms (DCT) presented in Sect. 4. A second contribution is a comparative evaluation of the different parameters (e.g. window size, feature representation and classifier) of our activity recognition system against several state-of-the-art benchmarks in Sect. 5.[2] The insights from the evaluation are designed to inform and serve as guidance for selecting effective parameter values when developing an activity recognition system. The data collection method introduced in this paper is also unique, in that it demon-

---

[2]Code and data associated with this paper are accessible from https://github.com/selfback/activity-recognition.
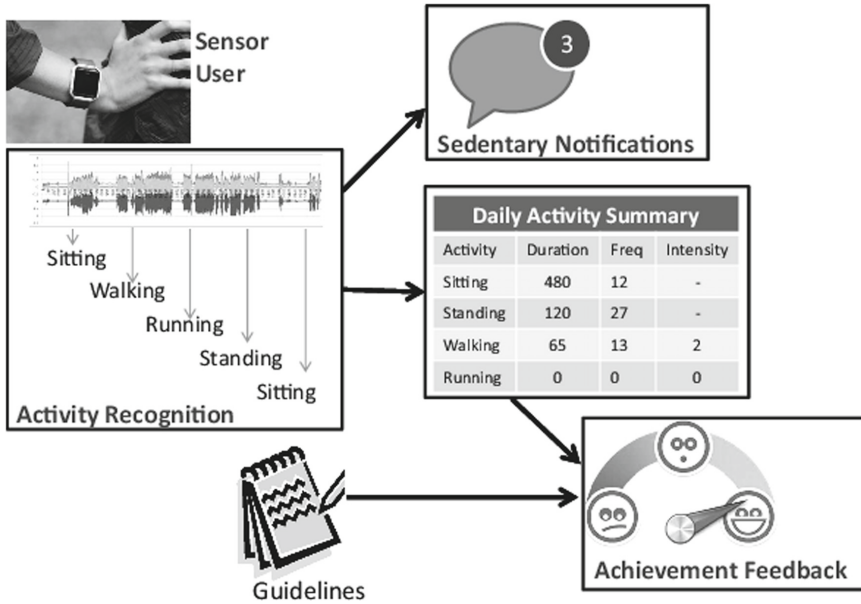
Fig. 1 Overview of SELFBACK system

strates how a script-driven method can be exploited to avoid the demand on manual transcription of sensor data streams (see Sect. 3). Related work and conclusions are also discussed and appear in Sects. 2 and 6 respectively.

2   Related Work in Activity Recognition

Physical activity recognition is receiving increasing interest in the areas of health care and fitness [13]. This is largely motivated by the need to find creative ways to encourage physical activity in order to combat the health implications of sedentary behaviour which is characteristic of today's population. Physical activity recognition is the computational discovery of human activity from sensor data. In the SELFBACK system, we focus on sensor input from a tri-axial accelerometer mounted on a person's wrist.

A tri-axial accelerometer sensor measures changes in acceleration in 3 dimensional space [13]. Other types of wearable sensors have also been proposed e.g. gyroscope. A recent study compared the use of accelerometer, gyroscope and magnetometer for activity recognition [17]. The study found the gyroscope alone was effective for activity recognition while the magnetometer alone was less useful. However, the accelerometer still produced the best activity recognition accuracy. Other sensors that have been used include heart rate monitor [18], light and temperature

sensors [16]. These sensors are however typically used in combination with the accelerometer rather than independently.

Some studies have proposed the use of a multiplicity of accelerometers [4, 15] or combination of accelerometer and other sensor types placed at different locations on the body. These configurations however have very limited practical use outside of a laboratory setting. In addition, limited improvements have been reported from using multiple sensors for recognising every day activities [9] which may not justify the inconvenience, especially as this may hinder the real-world adoption of the activity recognition system. For these reasons, some studies e.g. [14] have limited themselves to using single accelerometers which is also the case for SELFBACK.

Another important consideration is the placement of the sensor. Several body locations have been proposed e.g. thigh, hip, back, wrist and ankle. Many comparative studies exist that compare activity recognition performance at these different locations [4]. The wrist is considered the least intrusive location and has been shown to produce high accuracy especially for ambulation and upper-body activities [14]. Hence, this is the chosen sensor location for our system.

Many different feature extraction approaches have been proposed for accelerometer data for the purpose of activity recognition [13]. Most of these approaches involve extracting statistics e.g. mean, standard deviation, percentiles etc. on the raw accelerometer data (time domain features). Other works have shown frequency domain features extracted from applying Fast Fourier Transforms (FFT) to the raw data to be beneficial. Typically this requires a further preprocessing step applied to the resulting FFT coefficients in order to extract features that measure characteristics such as spectral energy, spectral entropy and dominant frequency [8]. Although both these approaches have produced good results, we use a novel approach that directly uses coefficients obtained from applying Discrete Cosine Transforms (DCT) on the raw accelerometer data as features. This is particularly attractive as it avoids further preprocessing of the data to extract features to generate instances for the classifiers.

3   Data Collection

Training data is required in order to train the activity recognition system. A group of 20 volunteer participants was used for data collection. All volunteers were either students or staff of Robert Gordon University. The age range of participants is 18 54 years and the gender distribution is 52 % Female and 48 % Male. Data collection concentrated on the activities provided in Table 1.

This set of activities was chosen because it represents the range of normal daily activities typically performed by most people. In addition, three different walking speeds (slow, normal and fast) were included in order to have an accurate estimate of the intensity of the activities performed by the user. Identifying intensity of activity is important because guidelines for health and well-being include recommendations for encouraging both moderate and vigorous physical activity [1].

Table 1 Details of activities used in our data collection script

| Activity name	Description
Walking slow	Walking at self-selected slow pace
Walking normal	Walking at self-selected normal pace
Walking fast	Walking at self-selected fast pace
Jogging	Jogging on a treadmill at self-selected speed
Up stairs	Walking up 4–6 flights of stairs
Down stairs	Walking down 4–6 a flights of stairs
Standing	Standing relatively still
Sitting	Sitting still with hands either on the desk or rested at the side
Lying	Lying down relatively still on a plinth
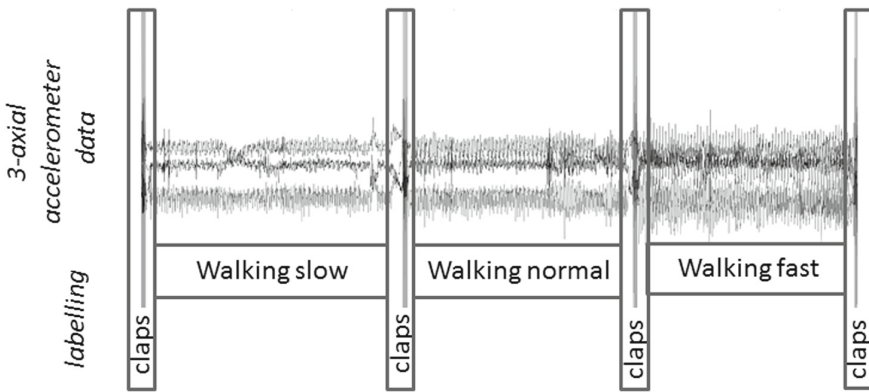
Fig. 2 Example of activity annotation with claps used to separate class transitions

Data was collected using the Axivity Ax3 tri-axial accelerometer[3] at a sampling rate of 100 Hz. Accelerometers were mounted on the wrists of participants using specially designed wristbands provided by Axivity. Participants were provided with scripts which contained related activities e.g. sitting and lying. The scripts guided participants on what activity they should do, how long they should spend on each activity (average of 3 min) and any specific details on how they should perform the activity e.g. sit with your arms on the desk.

Three claps are used to indicate the start and end of each activity. The three claps produce distinct spikes in the accelerometer signal which make it easy to detect the starts and ends of different activities in the data. This helps to simplify the annotation of the accelerometer data, by making it easy to isolate the sections of the data that correspond to specific activities. This allows the sections to be easily extracted and aligned with the correct activity labels from the script as shown in Fig. 2.

---

[3]http://axivity.com/product/ax3.

4 Activity Recognition Algorithm

The SELFBACK activity recognition system uses a supervised machine learning approach. This approach consists of 4 main steps which are: windowing, labelling, feature extraction and classifier training, as illustrated in Fig. 3.

4.1 Windowing

Windowing is the process of partitioning collected training data into smaller portions of length l, here specified in seconds. Figure 4 illustrates how windowing is applied to the 3-axis accelerometer data streams: x, y and z. Windows are overlapped by 0.5 of their length along the data stream. Thereafter each partitioned window, w, is used to generate an instance for activity classification. When choosing l, our goal is to find the window length that best balances between accuracy and latency. Shorter windows typically produce less accurate activity recognition performance, while longer windows produce latency, as several seconds worth of data need to be collected before a prediction is made. A comparative analysis of increasing window sizes ranging from 2 to 60 s is presented in Sect. 5.

4.2 Labelling

Once windows have been extracted, each window needs to be associated with a class label from a set of activity classes, $c \in C$. By default, this is the label of the activity

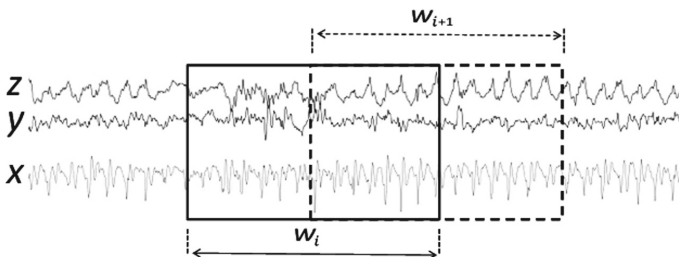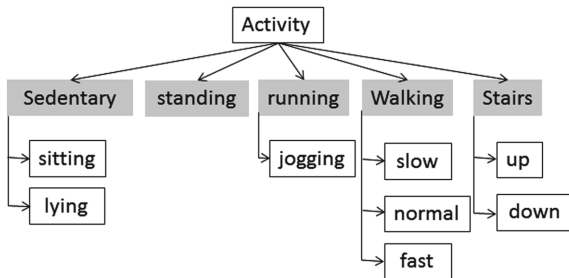**Fig. 3** SELFBACK activity recognition algorithm steps

Fig. 4 Illustration of accelerometer data windowing

Fig. 5 Activity class hierarchy

stream from which the window was extracted. Recall from Sect. 3 that $|C|$ was 9 classes (see Table 1), and can be thought of constituting a hierarchical structure as shown in Fig. 5. However, we observed that the more granular the activity labels, the more activity recognition accuracy suffers. In the case of some closely related classes e.g. sitting and lying, it is very difficult to distinguish between these classes from accelerometer data recorded from a wearable on the wrist. This is because wrist movement tends to be similar for these activities. Also, for activity classes distinguished by intensity (i.e. walking slow, walking normal and walking fast) the speed distinction between these activity classes can be more subjective than objective. Because the pace of walking is self-selected; one participant's slow walking pace might better match another's normal walking pace. Alternatively we consider $|C|$ equal to 5 classes by using the first level of the hierarchy (shaded nodes) with sub-tree raising of leaf nodes (whereby leaf nodes are grouped under their parent node). Evaluation results for activity recognition with both $|C|$ values are presented in Sect. 5.

4.3 Feature Extraction

The 3-axis accelerometer data streams, x, y and z, when partitioned according to the sliding window method as detailed in Sect. 4.1 generates a sequence of partitions, each of length l where each partition w_i is comprised of real-valued vectors x_i, y_i and z_i, such that $\mathbf{x} = (x_{i1}, \ldots, x_{il})$ DCT is applied to each axis (in essence each windowed partition x_i, y_i and z_i) to obtain a set of DCT coefficients which are an expression of the original accelerometer data in terms of a sum of cosine functions at different frequencies [10]. Accordingly the DCT-transformed vector representations, $\mathbf{x}' = \text{DCT}(\mathbf{x})$, $\mathbf{y}' = \text{DCT}(\mathbf{y})$ and $\mathbf{z}' = \text{DCT}(\mathbf{z})$, are obtained for each constituent in an instance. Additionally we derive a further magnitude vector, $\mathbf{m} = \{m_{i1}, \ldots, m_{il}\}$ of the accelerometer data for each instance as a separate axis, where m_{ij} is defined in Eq. 1.

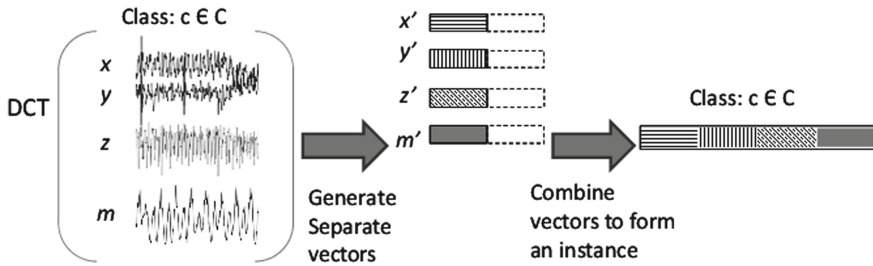$$m_{ij} = \sqrt{x_{ij}^2 + y_{ij}^2 + z_{ij}^2} \tag{1}$$

Fig. 6 Feature extraction and vector generation using DCT

As with $\mathbf{x}'$, $\mathbf{y}'$ and $\mathbf{z}'$, we also apply DCT to $\mathbf{m}$ to obtain $\mathbf{m}' = \text{DCT}(\mathbf{m})$. This means that our representation of a training instance consists of the pair $(\{\mathbf{x}', \mathbf{y}', \mathbf{z}', \mathbf{m}'\}, c)$, where c is the corresponding activity class label as detailed in Sect. 4.2. Including the magnitude in this way helps to train the classifier to be less sensitive to changes in orientation of the sensing device. Note that the coefficients returned after applying DCT are combinations of negative and positive real values. For the purpose of feature representation, we are only interested in the magnitude of the DCT coefficients, irrespective of (positive or negative) sign. Accordingly for each DCT coefficient e.g. x'_{ij}, we maintain its absolute value $|x'_{ij}|$.

DCT compresses all of the energy in the original data stream into as few coefficients as possible and returns an ordered sequence of coefficients such that the most significant information is concentrated at the lower indices of the sequence. This means that higher frequency DCT coefficients can be discarded without losing information. On the contrary, this might help to eliminate noise. Thus, in our approach we also retain a subset of the l coefficients and as proposed in [10] we retain the first 48 coefficients out of l. The final feature representation is obtained by concatenating the absolute values of the first 48 coefficients of $\mathbf{x}'$, $\mathbf{y}'$, $\mathbf{z}'$ and $\mathbf{m}'$ to produce a combined feature vector of length 192. An illustration of this feature selection and concatenation appears in Fig. 6.

4.4 Step Counting

An important piece of information that can be provided for ambulation activities is a count of the steps taken. This information has a number of valuable uses. Firstly, step counts provide a convenient goal for daily physical activity. Health research has suggested a daily step count of 10,000 steps for maintaining a desirable level of physical health [6]. A second benefit of step counting is that it provides an inexpensive method for estimating activity intensity. Step rate thresholds have been suggested in health literature that correspond to different activity intensities. For example, [1] identified that step counts of 94 and 125 steps per minute correspond to moderate and vigorous intensity activities respectively for men, and 99 and 135 steps per minute correspond to moderate and vigorous intensity activities for women. Accordingly,

step counts are likely to provide a more objective measure for activity intensity in the SELFBACK system than classifying different walking speeds. Here, we discuss two commonly used approaches involving frequency analysis and peak counting algorithms for inferring step counts from accelerometer data specific to ambulation activity classes.

4.4.1 Frequency Analysis

The main premise of this approach is that frequency analysis of walking data should reveal the heel strike frequency (i.e. the frequency with which the foot strikes the ground when walking) which should give an idea of the number of steps present in the data [2]. For walking data collected from a wrist-worn accelerometer, one or two dominant frequencies can be observed, heel strike frequency, which should always be present, and the arm swing frequency which may sometimes be absent. Converting accelerometer data from the time domain to the frequency domain using FFT enables the detection of these frequencies. For step counting, this approach seeks to isolate the heel strike frequency. Accordingly, the step count can be computed as a function of the heel strike frequency. For example, for frequency values in Hertz (cycles per second), the step count can be obtained by multiplying the identified heel strike frequency with the duration of the input data stream in seconds.

4.4.2 Peak Counting

The second approach involves counting peaks on low-pass filtered accelerometer data where each peak corresponds to a step. This process is illustrated in Fig. 7. For filtering, we use a Butterworth low-pass filter with a frequency threshold of 2 Hz for walking and 3 Hz for running.

The low-pass filter is then applied on m, which is the magnitude axis of the accelerometer signal obtained by combining the x, y and z axes. As a result, we expect to filter out all frequencies in m that are outside of the range for walking and running respectively. In this way, any changes in acceleration left in m can be attributed to the effect of walking or running. A peak counting algorithm is then deployed to count the peaks in m where the number of peaks directly corresponds to the count of steps.
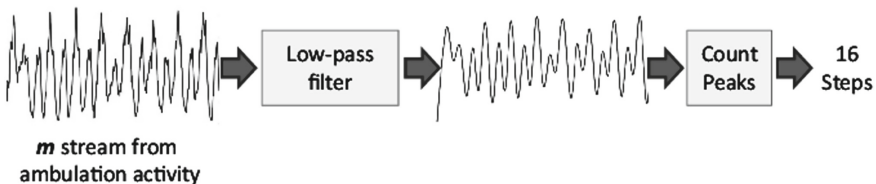

Fig. 7 Step counting using peak counting approach

5 Evaluation

In this section we present results for comparative studies that have guided the development of the SELFBACK activity recognition system. Firstly, an analysis of how window size and feature representation impact the effectiveness of human activity recognition is presented. Thereafter, we explore how classification granularity is affected by inter-class relationships and how that in turn impacts model learning. A question closely related with classification granularity is how to determine the activity intensity. For ambulation activities, step rate is a very useful heuristic for achieving this. Accordingly, we present comparative results for two step counting algorithms.

Our experiments are reported using a dataset of 20 users. Evaluations are conducted using a leave-one-person-out methodology i.e. one user is used for testing and the remaining 19 are used for training. In this way, we are testing the general applicability of the system to users whose data is not included in the trained model. Performance is reported using macro-averaged F1. SVM is used for classification after a comparative evaluation demonstrated its F1 score of 0.906 to be superior to that of kNN, decision tree, Nave Bayes and Logistic Regression; by more then 5 %, 12 %, 25 % and 3 % respectively.

5.1 Feature Representation and Window Size

For feature representation, we compare DCT, statistical time domain and FFT frequency domain features. Here time domain features are adopted from [19]. Figure 8 plots F1 scores for increasing window sizes from 2 to 60 s for each feature representation scheme.

The best F1 score is achieved with DCT features with a window size of 10 (F1 = 0.906). It is interesting to note that neither time or frequency domain features can match performance to that of directly using DCT coefficients for representation. Overall there is a 5 % gain in F1 scores with DCT compared to the best results of the rest.

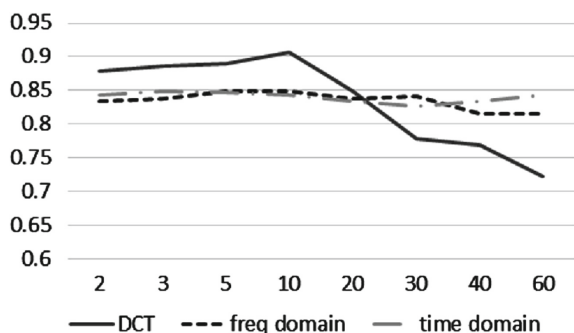**Fig. 8** Activity recognition performance at different window sizes

Table 2 Confusion matrix for 9-class activity classification

	Lying	Sitting	Standing	Jogging	Upstairs	Down stairs	Walk fast	Walk normal	Walk slow
Lying	115	127	8	0	0	0	0	0	1
Sitting	84	161	4	2	1	0	0	0	0
Standing	6	6	212	2	2	0	0	0	11
Jogging	0	0	0	284	1	0	0	1	0
Up stairs	0	0	3	8	92	7	7	11	30
Down stairs	1	0	1	5	31	89	5	4	8
Walk fast	0	0	1	21	3	1	157	53	6
Walk normal	0	0	1	4	11	3	48	141	41
Walk slow	0	1	7	3	23	4	0	32	181

5.2 Classification Granularity

Recall from Sect. 4.2 that data was collected relative to 9 different activities. Here we analyse classification accuracy with focus on inter-class relationships. In particular we study the separability of classes to establish which specific classes are best considered under a more general class of activity.

Overall F1 score for activity classification using 9 classes remains low at 0.688. Its confusion matrix is provided in Table 2, where the columns represent the predicted classes and the rows represent the actual classes. Close examination of the matrix shows that the main contributors to this low F1 score are due to classification errors involving activities lying, walking normal and upstairs. For instance we can see that for the activity class lying, only 115 instances are correctly classified and 125 instances are incorrectly classified as sitting. Similarly, 84 instances of sitting are incorrectly classified as lying. This indicates a greater discrimination confusion between lying and sitting which can be explained by wrist movement alone being insufficient to differentiate between these activities with a wrist worn accelerometer. However, both sitting and lying do represent sedentary behaviour and as such could naturally be categorised under the more general Sedentary class. A similar explanation follows for walking normal, where 48 instances are incorrectly classified as walking fast and 41 as walking slow. Accelerometer data for walking at different speeds will naturally be very similar. Also, the same walking speed is likely to be different between participants due to the subjectivity inherent in users judgment about their walking speeds. In addition, a user may unnaturally vary their pace while trying to adhere to a specific walking speed under data collection conditions. Again these reasons make it more useful to have the three walking speeds combined into one general class called Walking and have walking speed computed as a separate function of step rate. Regarding walking upstairs, we can see that it is most confused with walking slow but also suggests difficulties with differentiating between walking normal, walking fast and jogging. Many of these errors are likely to be addressed by

Table 3 Confusion matrix for 5-class activity classification

	Sedentary	Standing	Jogging	Stairs	Walking
Sedentary	490	7	2	1	3
Standing	17	205	2	1	14
Jogging	0	0	283	0	3
Stairs	3	0	9	0 223	67
Walking	3	5	24	31	679

taking into account inter-class relationships to form more general classes instead of having too many specialised classes.

Accordingly with the 5 class problem we have attempted to organise class membership under more general classes to avoid the inherent challenge of discriminating between specialised classes (e.g. between normal and fast walking). Therefore, there is a sedentary class combining sitting and lying classes; a stairs class to cover both upstairs and downstairs and a single walking class bringing together all different paces of walking speeds (See Fig. 5). Jogging and Standing remain as distinct classes as before.

As expected results in Table 3 shows that, 4 of the 5 classes have F1 scores greater than 0.9 with only Stairs achieving a score of 0.8. This result is far more acceptable than that achieved with the 9 class problem. The relatively lower F1 score with Stairs is due to 67 instances being incorrectly classified as Walking. This highlights the difficulty with differentiating between walking on a flat surface versus walking up or down stairs. However apart from the inclination of the surface there is no other characteristic that can help to differentiate these seemingly similar movements.

5.3 Step Counting

This final sub-section presents an evaluation of our step counting algorithms. For this, we collected a separate set of walking and running data with known actual step counts. This was necessary because actual counts of steps were not recorded for the initial dataset collected. In total, 19 data instances were collected for walking and 11 for running. For walking, participants were asked to walk up and down a corridor while counting the number of steps they took from start to finish. Reported step counts for walking range from 244 to 293. Participants performed a number of different hand positions which included walking with normal hand movement, with hands in trouser pocket and carrying a book or coffee mug. Walking data also included one instance of walking down a set of stairs (82 steps) and one instance of walking up a set of stairs (78 steps).

Table 4 Performance of step counting approaches measured using Root Mean Squared Error

Step counting approach	RMSE walking	RMSE running
Frequency analysis	11.245	6.250
Peak counting	**6.374**	**5.576**

Running data was collected on a treadmill. Participants were requested to run on a treadmill at a self-selected speed for a self-selected duration of time. Here also, three claps were used to mark the start and end of the running session. Two participants standing on the side were asked to count the steps in addition to the runner, due to the difficulty that may be involved in running and counting steps at the same time. Reported step counts for running range from 150 to 210.

The objective of this evaluation is to match, for each data instance, the count of steps predicted by each algorithm, to the actual step counts recorded. Root means squared error (RMSE) is used to measure performance. Because both step counting algorithms do not require any training, all 30 data instances are used for testing. Evaluation results are presented in Table 4. Generally it is useful to have mean squared error values that are below 10 for step counts. Overall we can see that better performance is observed from the Peak Counting method, thus this has been set as the default step counting approach for the SELFBACK system.

6 Conclusion

This paper focuses on the activity recognition part of the SELFBACK system which helps to monitor how well users are adhering to recommended daily physical activity for self-management of low back pain. The input into the activity recognition system is tri-axial accelerometer data from a wrist-worn sensor.

Activity recognition from the input is achieved using a supervised machine learning approach. This is composed of 4 stages: windowing, feature extraction, labelling and classifier training. Our results show that a window size of 10 s is best for identifying SELFBACK activity classes and highlighted the inherent challenge in differentiating between similar movement classes (such as lying with sitting and different paces of walking) using a wrist-worn sensor. Our approach to using Discrete Cosine Transform to represent instances achieved a 5 % classification performance gain over time and frequency domain feature representations. Algorithms to infer step counts from ambulation data suggests a simple peak counting approach following a low pass filter applied to the magnitude of the tri-axial data to be best. Future work will explore techniques for recognising a larger set of dynamically changing activities using incremental learning and semi-supervised approaches.

References

1. Abel, M., Hannon, J., Mullineaux, D., Beighle, A., et al.: Determination of step rate thresholds corresponding to physical activity intensity classifications in adults. J. Phys. Activity Health **8**(1), 45–51 (2011)
2. Ahanathapillai, V., Amor, J.D., Goodwin, Z., James, C.J.: Preliminary study on activity monitoring using an android smart-watch. Healthc. Technol. Lett. **2**(1), 34–39 (2015)
3. Airaksinen, O., Brox, J., Cedraschi, C.O., Hildebrandt, J., Klaber-Moffett, J., Kovacs, F., Mannion, A., Reis, S., Staal, J., Ursin, H., et al.: Chapter 4 european guidelines for the management of chronic nonspecific low back pain. Eur. Spine J. **15**, s192–s300 (2006)
4. Bao, L., Intille, S.S.: Activity recognition from user-annotated acceleration data. In: Pervasive Computing, pp. 1–17. Springer (2004)
5. Buchbinder, R., Blyth, F., March, L., Brooks, P., Woolf, A., Hoy, D.: Placing the global burden of low back pain in context. Best Pract. Res. Clin. Rheumatol. **27**, 575–589 (2013)
6. Choi, B.C., Pak, A.W., Choi, J.C.: Daily step goal of 10,000 steps: a literature review. Clin. Investig. Med. **30**(3), 146–151 (2007)
7. Dagenais, S., Caro, J., Haldeman, S.: A systematic review of low back pain cost of illness studies in the united states and internationally. Spine J. **8**(1), 8–20 (2008)
8. Figo, D., Diniz, P.C., Ferreira, D.R., Cardoso, J.M.: Preprocessing techniques for context recognition from accelerometer data. Pers. Ubiquit. Comput. **14**(7), 645–662 (2010)
9. Gao, L., Bourke, A., Nelson, J.: Evaluation of accelerometer based multi-sensor versus single-sensor activity recognition systems. Med. Eng. Phys. **36**(6), 779–785 (2014)
10. He, Z., Jin, L.: Activity recognition from acceleration data based on discrete consine transform and svm. In: Systems, Man and Cybernetics, 2009. SMC 2009. IEEE International Conference on, pp. 5041–5044. IEEE (2009)
11. Jordan, K.P., Kadam, U.T., Hayward, R., Porcheret, M., Young, C., Croft, P.: Annual consultation prevalence of regional musculoskeletal problems in primary care: an observational study. BMC Musculoskelet. Disord. **11**(1), 1 (2010)
12. Kent, P.M., Keating, J.L.: The epidemiology of low back pain in primary care. Chiropractic Osteopathy **13**(1), 1 (2005)
13. Lara, O.D., Labrador, M.A.: A survey on human activity recognition using wearable sensors. IEEE Commun. Surv. Tutorials **15**(3), 1192–1209 (2013)
14. Mannini, A., Intille, S.S., Rosenberger, M., Sabatini, A.M., Haskell, W.: Activity recognition using a single accelerometer placed at the wrist or ankle. Med. Sci. Sports Exerc. **45**(11), 2193 (2013)
15. Mäntyjärvi, J., Himberg, J., Seppänen, T.: Recognizing human motion with multiple acceleration sensors. In: 2001 IEEE International Conference on Systems, Man, and Cybernetics, vol. 2, pp. 747–752. IEEE (2001)
16. Maurer, U., Smailagic, A., Siewiorek, D.P., Deisher, M.: Activity recognition and monitoring using multiple sensors on different body positions. In: International Workshop on Wearable and Implantable Body Sensor Networks. BSN 2006, pp. 4–pp. IEEE (2006)
17. Shoaib, M., Bosch, S., Incel, O.D., Scholten, H., Havinga, P.J.: Fusion of smartphone motion sensors for physical activity recognition. Sensors **14**(6), 10146–10176 (2014)
18. Tapia, E.M., Intille, S.S., Haskell, W., Larson, K., Wright, J., King, A., Friedman, R.: Real-time recognition of physical activities and their intensities using wireless accelerometers and a heart rate monitor. In: 2007 11th IEEE International Symposium on Wearable Computers, pp. 37–40. IEEE (2007)
19. Zheng, Y., Wong, W.K., Guan, X., Trost, S.: Physical activity recognition from accelerometer data using a multi-scale ensemble method. In: IAAI (2013)