Comparative Study of Segmentation of Periodic Motion Data for Mobile Gait Analysis

Aris Valtazanos
School of Informatics
University of Edinburgh
Edinburgh, EH8 9AB
A.Valtazanos@sms.ed.ac.uk

D.K. Arvind
School of Informatics
University of Edinburgh
Edinburgh, EH8 9AB
dka@inf.ed.ac.uk

S. Ramamoorthy
School of Informatics
University of Edinburgh
Edinburgh, EH8 9AB
s.ramamoorthy@ed.ac.uk

ABSTRACT

Two approaches are presented and compared for segmenting motion data from on-body Orient wireless motion capture system for mobile gait analysis. The first is a ***basic, model-based*** algorithm which operates directly on the joint angles computed by the Orient sensor devices. The second is a ***model-free, Latent Space*** algorithm, which first *aggregates* all the sensor data, and then embeds them in a low-dimensional manifold to perform segmentation. The two approaches are compared for segmenting four different styles of walking, and then applied in a hospital-based clinical study for analysing the motion of elderly patients recovering from a fall.

Keywords

Wireless Sensor Networks, Motion Segmentation, Gait Analysis

1. INTRODUCTION

This paper is concerned with mobile gait analysis using a network of on-body, wireless Orient specks [24] for capturing the 3-D motion of the lower body. There is a need for recording the walking patterns of people for clinical study, and for objective methods for analysing them. At present the medical staff are limited to visually inspecting the patients' gait patterns. Gait laboratories with access to optical motion capture systems only record a snapshot of the gait on level, carpeted surfaces when the patient visits the clinic. It does not record the gait of the patient at other times of the day, or the gait on uneven surfaces or when climbing stairs or negotiating slopes. Moreover, the gait laboratories require highly trained staff to interpret the data, which leads to significant delays between the patient's interview, diagnosis and feedback.

The compact Orient device weighs 13gms and is contained in a perspex package measuring 36x28x11mm. With multiple Orient specks attached to the body parts, their measurements are synchronised and their results transmitted across the radio channel in sequence, so that the data of a complete frame can be assembled at the base-station within a few milliseconds. The data from three 3-axes accelerometer, gyroscope and magnetometer sensors is sampled at up to 512 Hz, and an orientation update rate of up to 64 Hz is achieved over the wireless network for full-body motion capture using a modest low-power 250 kbs radio [24].

The speck devices when positioned with their own internal axes along the traditional anatomical planes (sagittal, coronal, transverse), then the data output should be an authentic representation of the motion of the segment, irrespective of its precise location. This is in contrast with 3-D gait analysis using optical mo-cap which collects position information from surface markers which should be accurately positioned in relation to anatomical landmarks, which requires specific expertise.

2. RELATED WORK

Laptev *et al* [10] have identified and verified the existence of periodic patterns in optical tracking data from multiple cameras. Jean *et al* [9] employ a single camera, but are limited to motion tracking and do not address the more complex segmentation problem. Other methods include identifying local cyclic motions in body silhouettes extracted from video which are combined to obtain a global segmentation [2, 15], analyzing optical flow to identify temporal discontinuities in motion [18], and detecting periodic motions by using heuristics such as oscillations in motion intensity [21]. Blake *et al* [4] combine unsupervised learning with visual contour tracking in order to predict different classes of motions. Lv and Nevatia [12] divide the data points into feature sets which are segmented using Hidden Markov Models [16] and boosting techniques. All these methods are based on obtaining suitable optical data and its quality is compromised by noise, occlusion and limited visual fields.

A network of wireless inertial sensors on the person has been proposed as an alternative to optical systems for identifying periodic patterns in movement. A localised approach to segmentation and activity detection has been advocated which treats the individual nodes as independent entities [5, 8]. This method is similar to the model-based algorithm presented in Section 3.1. More advanced models of human walking have been used in [19, 14] for segmenting and determining spatio-temporal features of walking motion. In contrast, a simple kinematic model has been employed in the model-based algorithm in Section 3.1.

The concept of dimensionality reduction in segmentation is considered in [23] (see [13] for a related application for humanoid

Wireless Health '10, October 5–7, 2010, San Diego, USA

robots). They use the manifold primarily for action recognition, where each action class is modeled as a linear subspace model. Segmentation is performed by first creating a set of pre-segmented training examples, and using them as hypotheses in the context of online segmentation. This supervised approach is similar to the Hidden Markov Model-based segmentation [7, 6], which also requires the definition of prior transition probabilities, in order to model observation sequences. The segmentation is performed in a *distributed* fashion; in contrast, the Latent Space algorithm presented in Section 3.2 first *aggregates* all the sensor data, and then embeds them in a low-dimensional manifold to perform segmentation.

Our Latent Space approach is novel in taking an *unsupervised* approach to motion segmentation of inertial sensor data. The algorithm can be applied directly to arbitrary motion sequences, without the need to be trained on prior pre-segmented examples. The key to its functionality is the application of dimensionality reduction directly on each motion sequence. In effect the motion is summarised as a one-dimensional feature vector, which can be segmented using quite simple heuristics.

3. METHOD

We present and compare two periodic motion segmentation algorithms, which can also be used to detect intervals of motion activity in large motion sequences. The first is a ***basic, model-based*** algorithm which operates directly on the joint angles computed by the on-body Orient wireless motion capture network [24]. The second is a ***model-free, Latent Space*** algorithm, which takes as its input a *low-dimensional* representation of the human motion, and uses it to identify points and regions of interest.

3.1 Model-based algorithm

The model-based algorithm takes a greedy approach to motion segmentation, by attempting to identify local minima and maxima in periodic sequences. A key feature of this algorithm is the definition of a *human body model*, which reflects the body morphology of the tracked person.

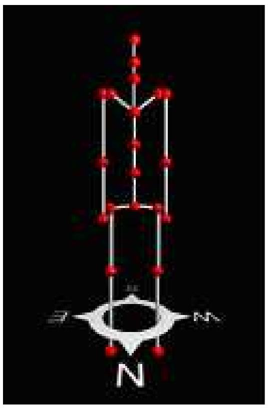

Figure 1: Representation of human body model.

The red dots in Figure 1 correspond to *joints* and the white links to *limbs*. The programmer measures the length of the various limbs, and passes them as inputs to the software program, so that the relative proportions of the tracked subject's limbs is correct. However, even if this measurement is not exact, it is still possible to use the relative lengths and positions of the joints to evaluate temporal features of the motion. Each node of the wireless sensor network is associated with a limb of interest in the body model. By aggregating the readings of the various sensors, the nodes compute the *angle* of the limbs connected to the respective joints [24].

A *frame* is defined as a collection of joint angles at a given moment in time. Joint angles are computed by aggregating information from the sensors (gyroscopes, accelerometers, magnetometers) in Orient, and combining them to obtain a single rotational estimate as described in [24]. For nodes of number N, a frame at time t is defined to be the set:

$$\{\theta_i\}_t, \quad i = 1..N \tag{1}$$

The position for all the *joints* in the body, joint positions, can be determined by applying forward kinematics by combining the computed joint angles and the limb lengths in the body model. At any time t, this is defined as the set:

$$\{\mathbf{p}_i\}_t, \quad i = 1..N, \quad \mathbf{p}_i = [x,\, y,\, z]^T \tag{2}$$

where x, y, z are the Cartesian coordinates of $\mathbf{p}_i$ relative to some *fixed* reference frame. This fixed point has been defined to be the hip joint of the model, so that $\mathbf{p}_{hip} = [0,\, 0,\, 0]^T$ for all times t. By considering the evolution of joint positions over time, a *motion sequence* of duration dt for joint i is defined as:

$$\{\{\mathbf{p}_i\}_t\}, \quad t \in [0, dt] \tag{3}$$

For every joint, this sequence is passed as input to the segmentation algorithm. The algorithm scans this sequence *greedily* for local minima and maxima in the joint positions. The three planes of motion are scanned independently, which results in three distinct sets per joint.

Let $j \in$ {transverse (x-z axis), sagittal (y-z axis), coronal (x-y axis)} denote the different planes of motion, and let $\{\mathbf{p}ij\}_t$ be the coordinate of the position of joint i on plane j at time t. Then, for every joint i and motion plane j, the *segmentation points* are a set of time instances:

$$\{\{\, t_{min}\}, \{\, t_{max}\}\} \tag{4}$$

such that every $\{\mathbf{p}_{ij}\}_{t_{min}}$ is a local minimum and $\{\mathbf{p}_{ij}\}_{t_{max}}$ is a local maximum. By considering pairs of successive minima and maxima, it is possible to identify distinct intervals in a periodic motion.

The model-based algorithm is fairly basic and is limited in several respects. Firstly, it is sensitive to sensory noise; slight local discontinuities in the joint positions will be treated as segmentation points, and results in *false positives*. Secondly, the model-based nature of the algorithm means that the designer must explicitly choose the most salient joints and planes of motion. This is not always an easy choice, even for simple periodic motion. For example, when a person is walking, is motion on the transverse plane more important than the one on the sagittal plane? Or, are the ankle positions more important in segmentation than the knee positions? Thirdly, even when such a choice has been made, it might still be possible to get conflicting segmentation sets between different joints. In such cases, deciding which joint produces more reliable results is a difficult task, and requires good modeling of sensor noise properties.

One may argue that thanks to these limitations, the basic algorithm lacks rigour in practice, while also reliant on the intervention and judgement of a healthcare or gait specialist.

3.2 Model-free algorithm

An alternative approach presented in this paper is an automated segmentation technique which can efficiently summarise periodicity. The *model-free, Latent Space* algorithm addresses these problems by removing the requirement to treat sensors and joints independently.

In this method the joint positions are grouped together into a single feature vector. The dimension of this vector, M, is three (the number of spatial dimensions) times the number of tracked joints, J. The net motion of a set of joints J is represented formally by the feature vector:

$$\mathbf{q} = [x_1,\, y_1,\, z_1,\, x_2,\, y_2,\, z_2, ..., x_t,\, y_t,\, z_t]^T \tag{5}$$

and, a *motion sequence* of duration *dt* can be now defined as a set of feature vectors:

$$\{\mathbf{q}_t\}, \quad t \in [0, dt] \tag{6}$$

Given the high dimensionality of the feature vectors, it would be difficult to apply the model-based segmentation algorithm to this motion sequence. However, the feature vectors can be thought of as originating from a *low-dimensional manifold*, where computational operations are more directly applicable. We refer to this manifold as the ***latent space*** of the motion. Therefore, given a high-dimensional motion sequence, our goal is to learn its latent space representation and use it to identify segmentation points.

There is a variety of manifold learning techniques available, most of which are used to find spectral embeddings of high-dimensional data points. Popular choices include the linear Principal Component Analysis [3] algorithm, and the non-linear Isomap [20], Locally Linear Embedding [17], and Gaussian Process Latent Variable Model [11] techniques. Non-linear manifold learning algorithms capture a wider range of dependencies and are therefore more suited to the motion sequences under analysis. Most of these algorithms take a set of high-dimensional points as input and compute a neighborhood graph for them. They subsequently use this graph to learn a latent space representation that successfully captures the neighborhood properties of the high-dimensional space.

For the purposes of our analysis, there are minor differences between the techniques, and we have chosen Isomap on the grounds of computational efficiency and speed. We have used the MATLAB Dimensionality Reduction toolbox developed in [22] for our experimental analysis.

By applying Isomap to a motion sequence $\{\mathbf{q}_t\}, \quad t \in [0, dt]$ with a target latent space dimension of 1, we obtain the *one - dimensional latent space motion sequence*:

$$\{\mathbf{r}_t\}, \quad t \in [0, dt] \tag{7}$$

The segmentation algorithm can now be applied to this motion sequence, as shown in Section 3.1. Thus, the new result will be a *single* set of segmentation points for the entire sequence, rather than several different sets for each of the joints and planes of motion. This model-free approach not only succeeds in simplifying segmentation, but in most cases also produces results comparable to the model-based one, as demonstrated in the Section 4.

3.3 Smoothing

Smoothing plays an important role in the model-based segmentation algorithm. The *smoothing factor* refers to the number of neighbouring points that are used to smooth a given point. If no smoothing is performed, then the model-based algorithm will tend to over-segment the motion. However, high smoothing is undesirable, because it tends to distort the “true" location of the segmentation points. By contrast, in the model-free algorithm, local anomalies are largely eliminated when the motion sequence is embedded in a lower-dimensional space. Thus, the smoothing factor may be reduced in this case significantly, and in most cases (unless explicitly stated so) we use a small smoothing factor of 5.

4. RESULTS

The segmentation algorithms were applied to a variety of motion datasets. Our experiments are divided into two sections. The first section deals with the analysis of walking motion, which exhibit clear periodic properties, whereas the second deals with the analysis of motion performed by elderly people aged over 85 years, who were recovering in the hospital after a fall. In all the cases, we present a comparison between the basic model-based algorithm, and the model-free, latent space segmentation algorithm. We show that the latent space algorithm can recover segmentations that closely match the corresponding points of the model-based algorithm. Supporting videos for all the datasets and results presented in this Section can be found in [1]. The videos also served as the primary means of visually validating the correctness of our results.

4.1 Evaluating periodic motions

In this set of experiments, the subject was asked to walk a few steps forward, turn around 180°, and walk forward for another few steps. The subject was asked to repeat this motion at various speeds and styles - normal, fast, slow, and shuffling walk.

4.1.1 Model-based segmentation

In this experiment, a normal walk with *six steps in each direction* is analysed. Figures 2 and 3 illustrate the difference between motion on the transverse and the sagittal plane, for the left and right knee joints. In Figure 3, we have also performed heavy smoothing

(smoothing factor = 40) and drawn the vertical lines to show the close correspondence between segmentations on the two planes. In each case, the position of the joint relative to the hips is plotted, with the positive direction being forward and upward. Note that although the unit is the same (centimeters) on both planes, the curves in each graph represent motion along different axes.

The *model-based* segmentations are presented in the form of red and black dots; red correspond to local minima and black to local maxima. Each segment bounded by a local minimum and a local maximum corresponds to one phase of the periodic motion. The knee was selected instead of the foot as it is only a limb away from the hips, which is used as a reference point. This choice avoids the redundancy that arises when joints are separated by multiple degrees of freedom.

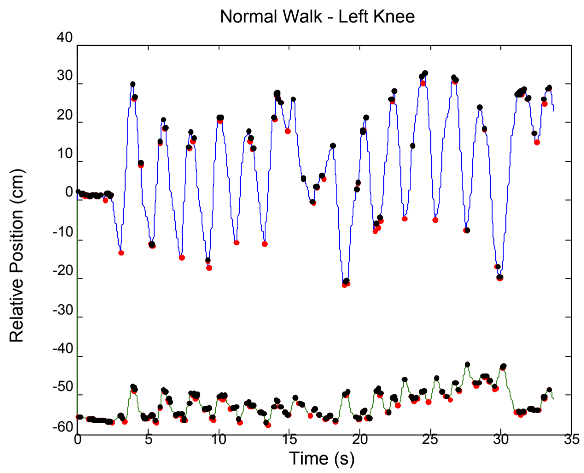

Figure 2: Left knee motion on transverse (top curve) and sagittal (bottom curve) plane (smoothing factor = 5)

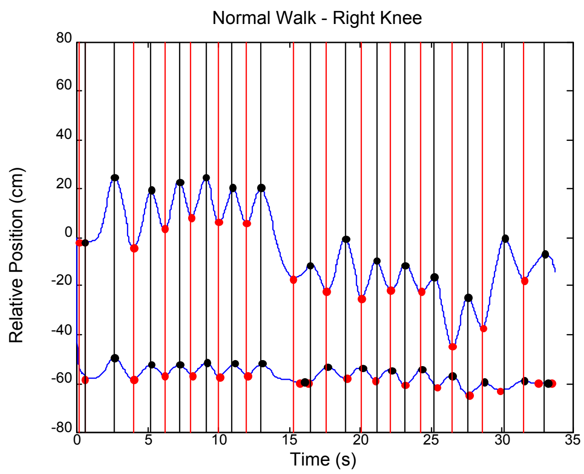

Figure 3: Overlaid segmentations for right knee motion, transverse (top curve) and sagittal (bottom curve) plane (smoothing factor = 40)

In the first half of the sequence, the minima and maxima on the two planes are closely aligned; the red and black lines go through the corresponding red and black points in both planes. The small gap just before time $t = 15$ indicates the 180° turn of the person. After that point, it can be seen that the segmentation points on each axis are out of phase, so a forward minimum occurs roughly at the same time as a lateral maximum and vice versa. The shift is explained by the 180° "reversal" of the coordinate system, which means that the foot is now extended towards the negative forward direction. This discrepancy could be avoided by normalising the coordinate system with respect to a fixed direction, for example, North. However, it would then not be possible to detect effects such as the 180° turn.

4.1.2 *Latent-space segmentation*

The identification of minima and maxima in the relative positions is an intuitive heuristic to characterise periodic motions such as walking. When a leg is fully extended during a walk, it is close to its furthest point from the waist, both in the forward and in the lateral direction - this is normally the point when the foot touches down. Equivalently, the minimum on the lateral direction occurs at the height of the foot swing phase; this will also be close to the point where the foot crosses the hip when advancing forward.

However, as explained in Section 3, the model-based algorithm can lead to mistakes if applied carelessly. Figure 2 is an example of this failure. Model-based segmentation correctly identifies all the extremes marking each step, but it also computes several false positives, as seen by the clutter of red and black points in the left-half of the figure. The number of false positives is also sensitive to smoothing, as the difference between Figures 2 and 3 demonstrates. Moreover, the choices between segmenting the knee or the ankle motion, or between segmenting the motion on the sagittal or on the transverse plane, is arbitrary and must be made by the human designer. As the previous section showed, it is also sometimes necessary to perform different types of comparison in order to verify segmentation results. For example, in Figure 3, we are cross-checking maxima and minima on the two planes to verify that they are correct; whereas if treated in isolation, it would be difficult to validate their correctness.

The latent space segmentation avoids these problems and moreover provides a more automated and heuristic-free way of characterising periodic motion. The method is compact as it summarises a multi-dimensional motion into a single dimension. In practice, it is a simple, fast, and effective way to identify segmentation points in periodic motion such as walking.

To illustrate the superiority of this approach, we reconsider the example of Figures 2-3, which suffered from the computation of several false positives. We first use the Isomap technique to reduce the multidimensional walking sequence to one dimension. Then, by considering the evolution of this one-dimensional trajectory over time, we apply the latent space segmentation algorithm on this latent space.

Figure 4 correctly recovers the segmentation points for the six steps in each direction, as well as the two steps of the turn. Note that the value plotted on the vertical axis is *unitless*, as it represents a spectral summary of the 12 dimensions of the original motion sequence (which correspond to three Cartesian space dimensions for each of the four joints - left/right, knee and ankle). In theory, the latent space represents values in centimeters, as each of the observed space dimensions represents coordinates in centimeters.

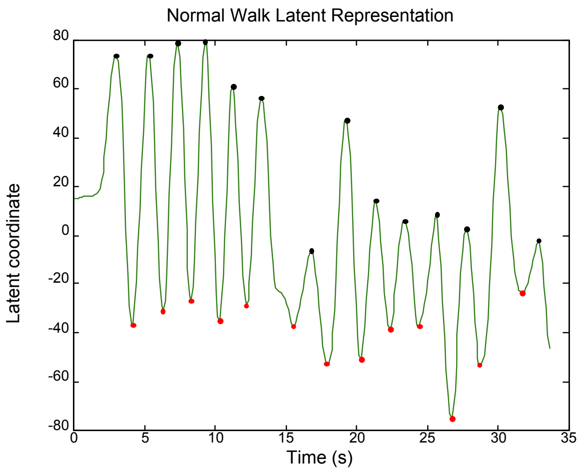

Figure 4: Latent space segmentation of the normal walking motion analysed in Figures 2-3

However, as previously mentioned, the dimensions span different planes, so when projected into a latent space, they lose their significance. Nevertheless, the temporal dimension remains unaltered, which allows one to estimate temporal features such as the period of the motion.

4.1.3 Application to different walking styles

This section discusses the application of the Latent Space algorithm to different walking styles. In all the cases, the subject repeated the same procedure as before; walk forward for a fixed distance, turn around by 180°, and walk forward for the same distance to end up at the original point. The difference in this case is that the number of steps taken to cover this distance varies, depending on the walking style used.

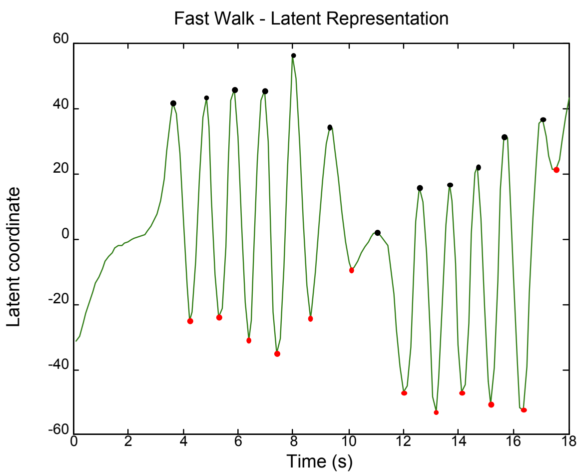

Figure 5: Fast walk - latent space segmentation

As Figures 5-7 demonstrate, the Latent Space algorithm manages to avoid the estimation of false positives. Furthermore, even though it was previously stated that the latent coordinate is unitless, it is still possible to make qualitative observations about the

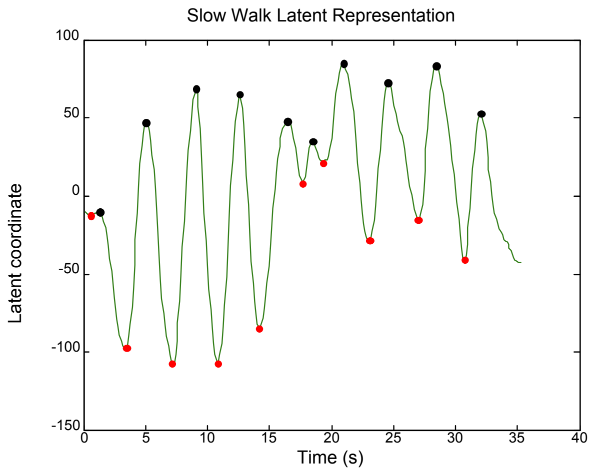

Figure 6: Slow walk - latent space segmentation

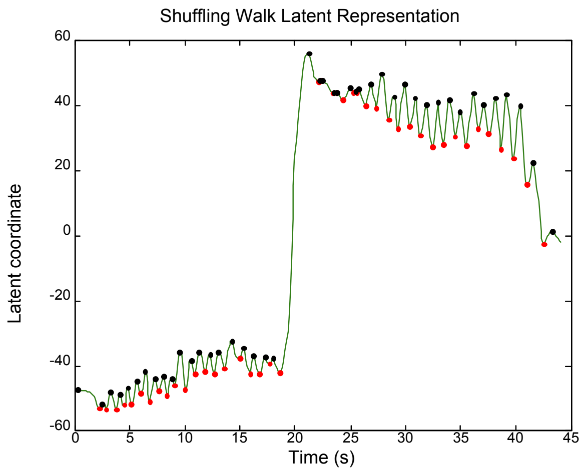

Figure 7: Shuffling walk - latent space segmentation

nature of the motion. This is most apparent in Figure 7, where the shuffling in the motion is analogous to the small variance in the latent coordinates. Intuitively, this small variance is a result of the small amount of movement that occurs in all three planes during a shuffling motion, which is directly propagated to the latent space representation. The turn that occurs halfway through the motion is also clearly distinguishable in this figure.

4.1.4 Walking period estimation

Table 1 summarises the results of the calculation of the average periodicity of the different walks, which accurately reflect the differences between the styles. For example, the period of the normal walk is smaller than the one of the slow walk but greater than the fast walk, as expected. Moreover, the computed period times are a strong indication that the motion is segmented at the correct granularity, without being either under- or over-segmented. The supporting videos [1] provide a means of visual validation and verification of these results.

	Average Period (s)
Fast walk	1.21
Normal walk	2.28
Slow walk	3.38
Shuffling walk	0.98

Table 1: Estimated mean walking period using latent space algorithm

4.1.5 Error estimation

The advantages of the Latent Space, model-free segmentation are evident on Figure 4. The algorithm correctly recovers all segmentation points that were identified by the model-based segmentation. Moreover, it avoids false positives by computing segmentations only at the boundaries of the six forward steps (and the two taken to turn around), in contrast to Figures 2-3 that over-segment the motion.

Nevertheless, it is important to represent any errors between the two methods in numerical terms. Figure 8 presents the error bars for the latent segmentation points computed in Figures 4-7. The *similarity error* margins are computed as follows: all segmentation points estimated from the positional information of all joints and planes of motion are collected into a single set (there were 12 subsets in this case, one for each dimension - see Section 4.1.2 for details). In both the high- and the low-dimensional space, the default smoothing factor of 5 has been used. Separate sets were formed for minima and maxima - we refer to the collection of all minima as *SPmins* and the collection of all maxima as *SPmaxs*. Then, for each latent segmentation maximum *lmax*, the similarity error is defined as:

$$\mathbf{err}(lmax) = \min_{m \in SPmaxs} |lmax - m| \tag{8}$$

and, correspondingly, for each latent minimum *lmin*:

$$\mathbf{err}(lmin) = \min_{m \in SPmins} |lmin - m| \tag{9}$$

Table 2 summarises the mean similarity error computed for each of the four latent motion sequences (fast, normal, slow, shuffling). For each sequence, the mean similarity error is compared to the mean walking period computed in Table 1, and the corresponding similarity error percentage is given.

	Mean Similarity Error (s)	Mean Walking Period (s)	Similarity Error Percentage
Fast walk	0.0417	1.21	3.45%
Normal walk	0.0392	2.28	1.72%
Slow walk	0.0471	3.38	1.39%
Shuffling walk	0.0352	0.98	3.59%

Table 2: Latent segmentation point mean similarity error, mean walking period and similarity error with respect to mean walking period

Considering that, in the latent space algorithm, a 12-dimensional

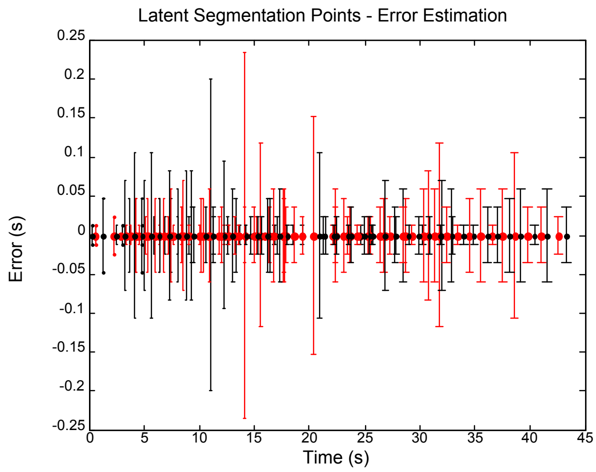

Figure 8: Similarity error estimation for the computed latent segmentation points of Figures 4-7 (red-minima, black-maxima). Note that, although the unit in both axes is seconds, the scales of the axes are *not* equal.

sequence has been effectively compressed into a single dimension, these error percentages are very small. The significance of this result is that for these experiments the latent space segmentation algorithm is successful in recovering and summarising the segmentation points computed at each individual joint and plane of motion, in the original high-dimensional space.

4.1.6 Performance with fewer sensor nodes

An additional advantage of the latent space algorithm is that it recovers segmentation points even when fewer sensor nodes are used. In the previous sections, 4 devices were used in total, leading to a total of 12 dimensions. In this section, we progressively reduce this number, and determine its effect on the performance of the latent space segmentation algorithm. We use the normal walk example of Figure 4 as an illustrative example.

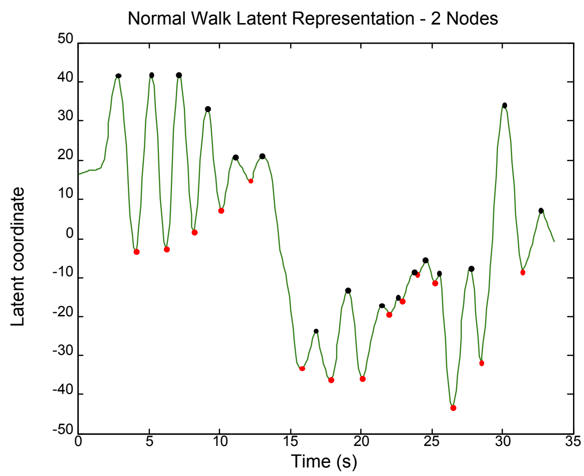

Figure 9: Normal walk - latent space representation using 2 sensor nodes (right/left knee)

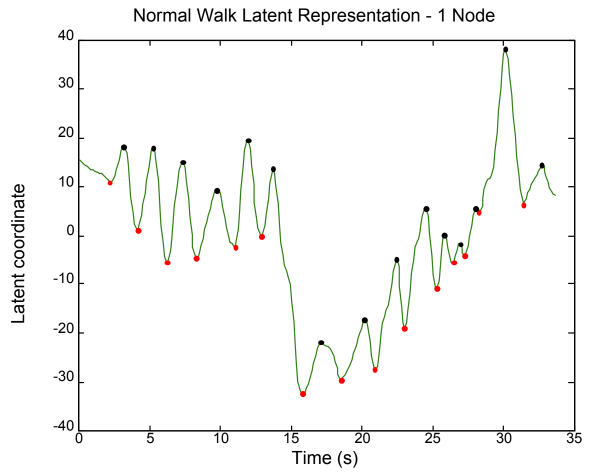

Figure 10: Normal walk - latent space representation using 1 sensor node (right knee)

In both figures, most of the segmentation points are recovered despite the reduced number of sensory data. The main weakness is the over-segmentation in the second half of both motion sequences, which is clearly inferior to the result of Figure 4. However, it is worth noting that the latent space algorithm in Figure 10 avoids the over-segmentations produced by the model-based variant of Figure 2, which also uses data from one sensor only and has a smoothing factor of 5.

4.2 Evaluating the motion of patients with mobility problems

4.2.1 Overview

The previous sections demonstrated how dimensionality reduction benefits segmentation of simple periodic motion sequences. However, the sequence of actions was known and thus interpretation of the segmented captures was possible. The full benefit of our approach would be seen in situations where the nature and variety of motions is not known precisely. An example of such an application is the analysis of the motion of elderly people who have suffered a fall. In such situations, the mobility of the subjects will be limited caused by the trauma and their gait may poorly resemble normal walking patterns. Thus, a more pertinent challenge in this case is identifying the intervals where waking occurs, and those intervals when the subjects are immobile. If this distinction can be made accurately, then segmentation would be a valuable tool in monitoring the rehabilitation of elderly people with mobility problems.

The datasets presented here were captured at the Royal Infirmary of Edinburgh. The subjects were elderly people over eighty years old who had recently suffered a fall, and were monitored during their recovery stages. Data was captured in two distinct phases. During the first phase, two patients (labelled Patient 1 and Patient 2) were asked to walk around a hospital room, occasionally taking a break to sit down or perform special exercises. The second phase involved the monitoring of a more seriously handicapped patient (Patient 3). As such, this subject was observed executing whatever motions he felt comfortable with, rather than being asked to walk or do specific kinds of exercise.

For these experiments, the *segmentation algorithms* have been transformed into *activity detection algorithms*. Rather than scanning for local minima and maxima, the algorithms now only segment at points that match a predefined movement intensity threshold. This threshold has been defined in order to distinguish between possible regions of activity and inactivity.

4.2.2 Model-based activity detection

As in the previous section, we begin with two model-based activity detection examples, for Patients 1 and 2. The algorithm scans for high intensity points on the sagittal plane, and overlays the computed values to the motion on the transverse plane - left ankle motion is plotted as an example.

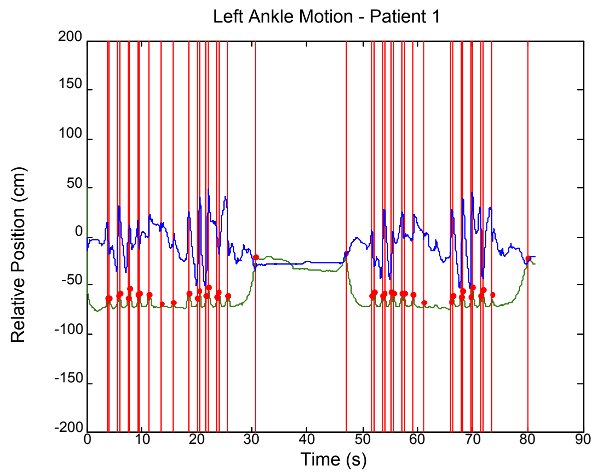

Figure 11: Patient 1 left ankle motion - transverse (top blue curve) and sagittal plane (bottom green curve)

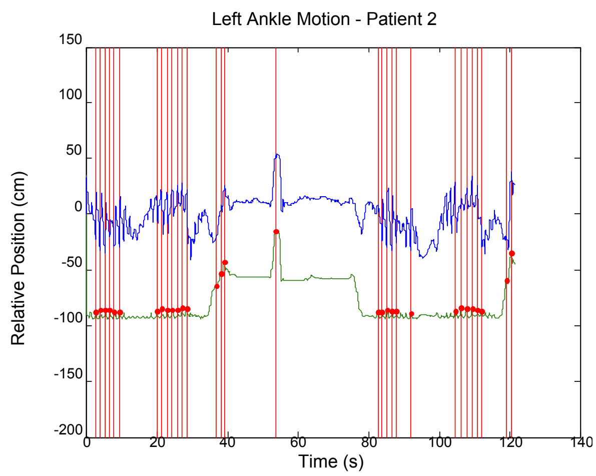

Figure 12: Patient 2 left ankle motion - transverse (top blue curve) and sagittal plane (bottom green curve)

Using this modified activity detection procedure, it is possible to identify clusters where walking motion occurs - these clusters are clearer for Patient 2. It can be seen that the segmentation points on the sagittal direction closely follow the regions of high motion intensity on the transverse direction - this correlation is particularly strong in Figure 12. This correspondence is a strong indication that the patient is performing some form of translational motion, such as walking.

Furthermore, the plateaus in the middle of all four figures mark regions where motion is not detected in either direction. The plateau on the sagittal direction is slightly elevated which implies that the distance between the feet and the hips has been reduced. Thus, the patient is performing some kind of bending motion, an example of which would be to sit down. Similarly, the spike in the middle of Figure 12 indicates that Patient 2 could be flexing his feet while sitting down. However, in a different context, such fluctuations could also be indicative of a fall or other irregular behaviour. The incorporation of additional data, such as gyroscope reading, in order to formally classify activity in such situations, can be treated as an extension to this work.

4.2.3 Latent space activity detection

In this context, it is more important to identify regions of activity and inactivity, rather than finding exact quantitative properties of segmented motion. Thus, the main expectation from a latent space activity detection algorithm is to give an immediate impression of where these regions occur, while avoiding the need for cross-validation between dimensions that arises with the basic model-based algorithm.

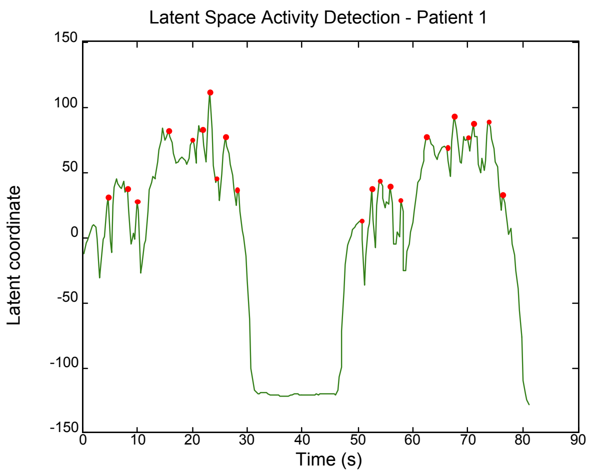

Figure 13: Latent space activity detection - Patient 1

Figures 13 and 14 demonstrate the application of the latent space, model-free activity detection algorithm, for the two motion sequences evaluated in Section 4.2.2. The red points signal moments where activity above the defined threshold has been detected. The occurrences of the plateaus are correctly recovered in both cases, as well as the spike just before $t = 60$ in Figure 12. Moreover, it is now possible to define four distinct clusters of intense motion in Figure 14, which were not as clear in Figures 11 and 12.

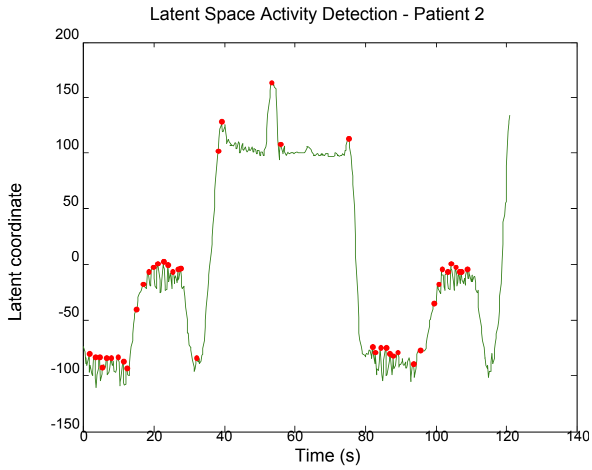

Figure 14: Latent space activity detection - Patient 2

4.2.4 Walking period estimation

Since the walking clusters can be easily identified in the latent space, it is possible to use them to estimate the period of the patient's walking motion. Table 3 summarises the results for each of the four walking clusters identified for patients 1 and 2 in the previous section. The period is estimated by performing *k-means clustering* on the detected high-activity points, and then calculating the period of each cluster. The results demonstrate a similarity in the walking period of the patients, for each of the identified intervals.

	Patient 1	**Patient 2**
Interval 1	1.52	1.06
Interval 2	1.53	1.19
Interval 3	1.44	1.02
Interval 4	1.39	1.01

Table 3: Estimated mean walking period, in seconds, for each of the four clusters identified in Figures 13 and 14

4.2.5 Activity detection for Patient 3

The activity detection algorithms were also applied to data from Patient 3, who was diagnosed with severer mobility problems.

Figure 15 indicates a different behaviour to the one identified in Figures 11-12. Because of the patient's mobility problems, it is now more difficult to discern intervals of activity. This feature is reflected in Figure 15. Although significant motion activity is observed on the transverse direction, hardly any can be seen on the sagittal plane. A possible explanation could be that the patient is performing some kind of shuffling motion, with his feet barely raised off the ground. However, regardless of the exact nature of the motion, it is essential that motion activity be easily identifiable, without the need to cross-validate data from different sensors and joints. Figure 16 shows the latent space representation of the above sequence, which effectively summarises the region where motion occurs in Figure 15.

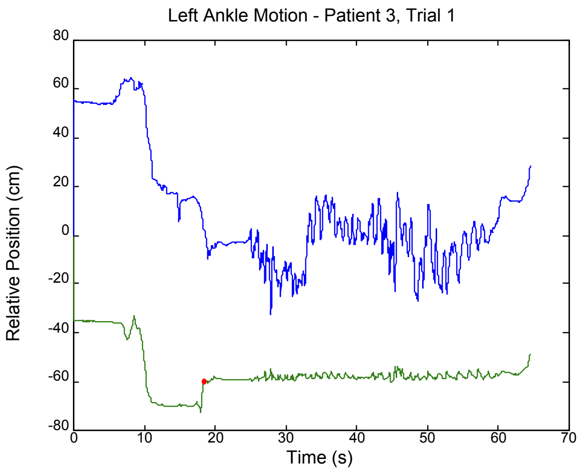

Figure 15: Patient 3 left ankle motion - transverse (top blue curve) and sagittal plane (bottom green curve)

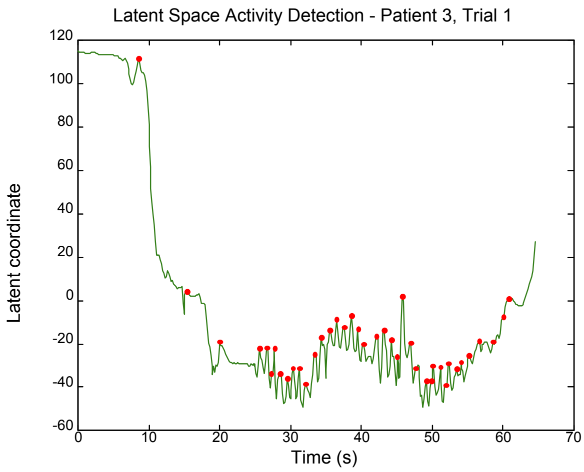

Figure 16: Latent space activity detection - Patient 3

5. CONCLUSION

Two algorithmic approaches have been presented for analysing a person's gait using the on-body wireless Orient motion capture system. The two algorithms are compared for the analysis of different styles of walking and then applied in a clinical study involving elderly patients recovering after a fall. Though both model-based and model-free activity detection can be useful in this context, the latter appears to be far more powerful in clustering data more strongly, and providing quantitative measures such as the walking period. Latent space summarisation techniques could be of great benefit in the monitoring of patients with mobility problems, during their recovery in the hospital which can continue in their homes after their discharge.

The main computational bottleneck in the application of the Latent Space algorithm is the estimation of the low-dimensional manifold, from a high-dimensional motion sequence. The limitation of manifold learning algorithms is that their complexity scales exponentially with the dataset size, and are efficient for sizes up to 700 points. For a set of this size, a one-dimensional manifold can be learnt in less than half a second; if this limit is exceeded, the process may take up to a few seconds or even tens of seconds for large datasets (close to 5,000 points).

This was a problem that had to be considered during the implementation given that many of the motion sequences - especially those involving patients - were lengthy. A simple yet effective solution proved to be the *pruning* of our datasets. Since the Orient sensors could provide orientation data at a frequency exceeding 100Hz, more data was available than was necessary to compute segmentation points. As such, it was possible to prune the datasets by a factor of 10 or even 20, without losing much crucial information.

However, even with the reduction of the data set it is still not possible to apply the Latent Space segmentation online, in its current form. The reason is the algorithmic procedure of manifold learning, requires that *neighborhood graphs* between data points are formed. In order to create an accurate neighborhood graph that reflects the topology of the space, all points must have been gathered; in this case, this means that the whole motion must be completed before applying segmentation. It would be worth investigating the learning of a mapping from the observed high-dimensional space to the latent space. This mapping would be learnt in a supervised manner and used to embed points to the latent space online. Nonetheless, learning such mappings that generalize well to unseen data points is a difficult task, and one that few existing manifold learning techniques address adequately.

A further limitation of the latent space algorithm is that it is restricted to *unimodal* periodic motion sequences. For example, in all the experiments, the relative positions of the knees and ankles were gathered and analysed jointly, under the assumption that they all affect the same, lower-body, periodic walking motion. However, if sensors were also placed on the upper body, with the aim of capturing the periodicity of hand motion during walking, the Latent Space algorithm would not be able to distinguish between upper- and lower-body periodicity. A solution to this problem would be to first cluster sensors into groups, depending on which periodic motion they affect, and then apply the algorithm to each group of sensors separately.

Although we demonstrated that segmentation can be performed without supervision, it would be difficult to use the same argument in action classification. In that context, it is necessary to exploit the qualitative and quantitative information provided by data, in order to distinguish between different classes of actions. Towards this end, future work will attempt to address this problem, by using the computed segments to categorize motions. Nevertheless, the main motivation behind our proposed approach is that good quality motion segmentation is a prerequisite for accurate and reliable motion recognition and classification.

6. ACKNOWLEDGMENTS

AV is supported by a doctoral studentship from the Research Consortium in Speckled Computing which is funded by the Scottish Funding Council as part of the Strategic Research Development Programme (R32329), and the UK Engineering and Physical Sciences Research Council (EPSRC) as part of the Basic Technology Programme (C523881).

7. REFERENCES

[1] Supporting material, http://www.specknet.org/publications.

[2] A. B. Albu, R. Bergevin, and S. Quirion. Generic temporal segmentation of cyclic human motion. *Pattern Recogn.*, 41(1):6–21, 2008.

[3] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, 1st ed. 2006. corr. 2nd printing edition, October 2007.

[4] A. Blake, M. Isard, and D. Reynard. Learning to track the visual motion of contours. *Artificial Intelligence*, 78:101–134, 1995.

[5] J. Chen, K. Kwong, D. Chang, J. Luk, and R. Bajcsy. Wearable sensors for reliable fall detection. In *Proceedings of the 2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*, 2005.

[6] E. Guenterberg, H. Ghasemzadeh, and R. Jafari. A distributed hidden markov model for fine-grained annotation in body sensor networks. In *BSN '09: Proceedings of the 2009 Sixth International Workshop on Wearable and Implantable Body Sensor Networks*, pages 339–344, Washington, DC, USA, 2009. IEEE Computer Society.

[7] E. Guenterberg, H. Ghasemzadeh, V. Loseu, and R. Jafari. Distributed continuous action recognition using a hidden markov model in body sensor networks. In *Proceedings of the 5th IEEE International Conference on Distributed Computing in Sensor Systems*, pages 145–158, Berlin, Heidelberg, 2009. Springer-Verlag.

[8] E. Guenterberg, S. Ostadabbas, H. Ghasemzadeh, and R. Jafari. An automatic segmentation technique in body sensor networks based on signal energy. In *BodyNets '09: Proceedings of the Fourth International Conference on Body Area Networks*, pages 1–7, ICST, Brussels, Belgium, Belgium, 2009. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

[9] F. Jean, R. Bergevin, and A. Branzan-Albu. Body tracking in human walk from monocular video sequences. In *Proceedings of the Second Canadian Conference on Computer and Robot Vision (CRV)*, pages 144–151, Victoria, BC, Canada, May 9-11 2005.

[10] I. Laptev, S. J. Belongie, P. Pérez, and J. Wills. Periodic motion detection and segmentation via approximate sequence alignment. In *In Proceedings of the International Conference on Computer Vision*, 2005.

[11] N. Lawrence. Gaussian process latent variable models for visualisation of high dimensional data. In *NIPS*, 2003.

[12] F. Lv and R. Nevatia. Recognition and segmentation of 3-d human action using hmm and multi-class adaboost. In A. Leonardis, H. Bischof, and A. Pinz, editors, *ECCV 2006*, volume 3954 of *Lecture Notes in Computer Science*, pages 359–372. Springer Berlin / Heidelberg, 2006.

[13] K. F. MacDorman, R. Chalodhorn, and M. Asada. Periodic nonlinear principal component neural networks for humanoid motion segmentation, generalization, and generation. *Pattern Recognition, International Conference on*, 4:537–540, 2004.

[14] S. Miyazaki. Long-term unrestrained measurement of stride length and walking velocity utilizing a piezoelectric gyroscope. In *IEEE Transactions on Biomedical Engineering*, volume 44, pages 753–759, 1997.

[15] S. Quirion, A. Branzan-albu, and R. Bergevin. Skeleton-based temporal segmentation of human activities. In *Proceedings of the 13th International Conference in Central Europe on Computer Graphics (WSCG 05)*, pages 145–148, 2005.

[16] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE*, pages 257–286, 1989.

[17] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, December 2000.

[18] Y. Rui and P. An. Segmenting visual actions based on spatio-temporal motion patterns. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2000.

[19] A. Sabatini, C. Martelloni, S. Scapellato, and F. Cavallo. Assessment of walking features from foot inertial sensing. In *IEEE Transactions on Biomedical Engineering*, volume 52, pages 486–494, 2005.

[20] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, Dec 2000.

[21] A. Thangali and S. Sclaroff. Periodic motion detection and estimation via space-time sampling. *Applications of Computer Vision and the IEEE Workshop on Motion and Video Computing, IEEE Workshop on*, 2:176–182, 2005.

[22] L. van der Maaten, E. Postma, and H. van den Herik. Dimensionality reduction: A comparative review. Technical Report TiCC-TR 2009-005), Tilburg University, 2009.

[23] A. Y. Yang, S. Iyengar, S. Sastry, R. Bajcsy, P. Kuryloski, and R. Jafari. Distributed segmentation and classification of human actions using a wearable sensor network. In *in Proceedings of the CVPR Workshop on Human Communicative Behavior Analysis*, 2008.

[24] A. D. Young, M. J. Ling, and D. K. Arvind. Orient-2: a realtime wireless posture tracking system using local orientation estimation. In *EmNets '07: Proceedings of the 4th workshop on Embedded networked sensors*, pages 53–57, New York, NY, USA, 2007. ACM.