

Latent Space Segmentation for Mobile Gait Analysis¹

ARIS VALTAZANOS and D.K. ARVIND and SUBRAMANIAN RAMAMOORTHY
School of Informatics, University of Edinburgh

An unsupervised learning algorithm is presented for segmentation and evaluation of motion data from the on-body Orient wireless motion capture system for mobile gait analysis. The algorithm is *model-free* and operates on the *latent space* of the motion, by first *aggregating* all the sensor data into a single vector, and then modeling them on a low-dimensional manifold to perform segmentation. The proposed approach is contrasted to a *basic, model-based* algorithm, which operates directly on the joint angles computed by the Orient sensor devices. The latent space algorithm is shown to be capable of retrieving qualitative features of the motion even in the face of noisy or incomplete sensor readings.

Categories and Subject Descriptors: H.4 [Information Systems and Applications]: Miscellaneous

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: Wireless Sensor Networks, Motion Segmentation, Gait Analysis

1. INTRODUCTION

There is a growing need for recording the walking patterns of people for clinical study, and for objective methods for analysing them. Automated analysis of motion data is hard, especially when no prior training examples are provided. At present medical staff are limited to visually inspecting patients' gait patterns. Gait laboratories with access to optical motion capture systems only record a snapshot of gait on level, carpeted surfaces when a patient visits the clinic. It does not record the gait of the patient at other times of the day, on uneven surfaces, when climbing stairs or when negotiating slopes. Moreover, gait laboratories require highly trained staff to interpret the data, which leads to significant delays between the patient's interview, diagnosis and feedback.

This paper is concerned with mobile gait analysis using a network of on-body, wireless Orient specks [Young et al. 2007] for capturing the 3-D motion of the lower body. 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

¹This paper is an extended journal version of the conference paper: A. Valtazanos, D.K. Arvind, S. Ramamoorthy, *Comparative study of segmentation of periodic motion data for mobile gait analysis*. In Proc. Wireless Health Conference, 2010.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 1529-3785/20YY/0700-0001 \$5.00

are synchronised and 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 [Young et al. 2007].

When positioned with their own internal axes along the traditional anatomical planes (sagittal, coronal, transverse), the speck devices provide 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 motion capture, which collects position information from surface markers. In such a context, the markers must be accurately positioned in relation to anatomical landmarks, which requires specific expertise and guidance.

In a previous study [Valtazanios et al. 2010], we presented results on unsupervised latent space gait segmentation for healthy subjects, and activity detection for elderly patients recovering from a fall. In this paper, we provide an extended version of the latent space algorithm and its properties, with additional patient results where sensory data was found to be noisy or incomplete. We show that latent space representation can correctly recover qualitative features of the motion even under these more restrictive conditions.

2. RELATED WORK

Laptev *et al* [Laptev et al. 2005] identify and verify the existence of periodic patterns in optical tracking data from multiple cameras. Jean *et al* [Jean et al. 2005] 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 [Albu et al. 2008][Quirion et al. 2005], analyzing optical flow to identify temporal discontinuities in motion [Rui and An 2000], and detecting periodic motions by using heuristics such as oscillations in motion intensity [Thangali and Sclaroff 2005]. Blake *et al* [Blake et al. 1995] combine unsupervised learning with visual contour tracking in order to predict different classes of motions. Lv and Nevatia [Lv and Nevatia 2006] divide the data points into feature sets which are segmented using Hidden Markov Models [Rabiner 1989] 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. Localised approaches to segmentation and activity detection treat individual nodes as independent entities [Chen et al. 2005][Guenterberg et al. 2009]. These methods are similar to the model-based algorithm presented in Section 3.1. More advanced models of human walking have been used in [Sabatini et al. 2005][Miyazaki 1997] 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.

Dimensionality reduction in segmentation is considered in [Yang et al. 2008] (see [MacDorman et al. 2004] for a related application for humanoid 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 [Guenterberg et al.

2009] [Guenterberg et al. 2009], which also requires the definition of prior transition probabilities to model observation sequences. The segmentation is performed in a *distributed* fashion; in contrast, our algorithm first collects all the sensor data into a single feature vector, and then embeds them in a low-dimensional manifold to perform segmentation.

Our latent space algorithm takes an *unsupervised* approach to motion segmentation of inertial sensor data. The advantage of this approach is that the algorithm can be applied directly to arbitrary novel motion sequences, without the need to be trained on prior pre-segmented examples. In effect, the algorithm summarises the motion as a one-dimensional feature vector, which can be segmented using simple heuristics.

The main focus of our paper is the precise characterisation of discontinuous events in periodic motions, such as walking. In particular, we are interested in extracting and describing the qualitative structure of motion sequences, whose exact features may be unknown. By taking a latent space approach to motion segmentation, we seek to characterise sequences in terms of their *critical points* (e.g. maxima, minima) in that space, and to illustrate how these points yield a concise representation of the evaluated task. Thus, our approach is largely experimental, focusing on the analysis of real clinical results from both patients and healthy subjects.

3. METHOD

We first present a basic, model-based algorithm, which serves as a benchmark for our analysis. This routine operates directly on the joint angles computed by the on-body Orient wireless motion capture network [Young et al. 2007]. We then present the main contribution of our paper, which is a *model-free, latent space* algorithm. This procedure takes as its input a *low-dimensional* representation of a motion sequence, and uses it to identify points and regions of interest. The algorithms are presented in the context of motion segmentation, though we later show how their applicability can be extended, with minor modifications, to activity detection and motion recovery under noisy sensing.

3.1 Model-based algorithm

The model-based algorithm takes a greedy approach to motion segmentation, by searching for local minima and maxima in periodic sequences. This algorithm relies on the structure of the human body model, which reflects the body morphology of the tracked person.

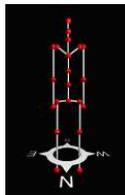


Fig. 1. Representation of human body model.

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 are 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. Thus, in principle, the same body model can be used regardless of the physical characteristics of the tracked person. Each node of the wireless sensor network is associated with a limb of interest in the body model. By joining the readings of the various sensors into a feature vector, the nodes compute the *angle* of the limbs connected to the respective joints [Young et al. 2007].

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 [Young et al. 2007]. For a total of N nodes, a frame at time t is the set:

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

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

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

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:

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

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, leading to three distinct sets per joint.

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

$$\{\{t_{min}\}, \{t_{max}\}\} \quad (4)$$

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. First, it is sensitive to sensory noise; slight local discontinuities in the joint positions will be treated as segmentation points, which leads to several *false positives*. Second, 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? Third, 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 one that requires good modeling of sensor noise properties.

3.2 Model-free algorithm

The main contribution of this paper is an alternative, 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 by the feature vector:

$$\mathbf{q} = [x_1, y_1, z_1, x_2, y_2, z_2, \dots, x_J, y_J, z_J]^T \quad (5)$$

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

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

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 points can be thought of as being drawn from a *low-dimensional manifold*, where computational operations such as segmentation are better posed. 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 [Bishop 2007] algorithm, and the non-linear Isometric Manifold Learning (Isomap) [Tenenbaum et al. 2000], Landmark Isomap [De Silva and Tenenbaum 2003], Locally Linear Embedding [Roweis and Saul 2000], and Gaussian Process Latent Variable Model [Lawrence 2003] techniques. Approximation algorithms such as Fast Isomap [Lei et al. 2010] have also been proposed recently, in an attempt to minimise the complexity of learning the structure of a high-dimensional dataset. Non-linear manifold learning algorithms capture a wider range of dependencies and are therefore more suited to the motion sequences under analysis. There exist several software implementations for dimensionality reduction algorithms; in our experiments, we use the MATLAB toolbox developed in [van der Maaten et al. 2009].

3.2.1 The Isomap algorithm. 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 briefly summarise the key features of this algorithm.

As with all manifold learning algorithms, the goal of Isomap is to find a suitable d -dimensional representation for a given D -dimensional dataset, where $d < D$. The value of d is often specified by the user and corresponds to the dimensionality of the latent space.

The algorithm first constructs a *neighbourhood graph* G for the high dimensional dataset, in order to capture the relationship between the various points in the set. A point i is connected in G if j is one of its k nearest neighbours, where k is also user-defined. If such a connection is formed, the length of the edge is equal to the Euclidean distance $d(i, j)$ between points i and j .

The neighbourhood graph is used to construct a matrix M holding the *shortest distances* between all pairs of points. Dijkstra's algorithm is a common choice for finding shortest paths in graph structures. Initially, the shortest path distance between two points i and j is set to $d(i, j)$ if an edge between them exists in G , and ∞ otherwise. The algorithm then gradually iterates over the k neighbours of each point, calculating the distance between each point-neighbour pair and updating shortest path distances accordingly.

The final step of the procedure is the d -dimensional embedding of the original dataset. This is achieved by obtaining, in descending order, the eigenvectors of the matrix

$$K = \frac{1}{2} H M^2 H \quad (7)$$

where $M_{ij}^2 = (M_{ij})^2$ (that is, M^2 is the elementwise square of the matrix M), and $H_{ij} = \delta_{ij}^2 - \frac{1}{N}$. The sorted eigenvectors and eigenvalues are used to obtain the latent coordinates of each embedded datapoint. If λ_p is the p -th eigenvalue and v_i^p is the i -th component of the p -th eigenvector, then the p -th component of the embedded point \mathbf{r} is equal to

$$\mathbf{r}^p = \sqrt{\lambda_p} \cdot v_i^p \quad (8)$$

In summary, Isomap first attempts to learn the structure of the high dimensional points through their neighbourhood graph, and then exploits this information to obtain their low-dimensional embedding. The definition of the low-dimensional coordinates in terms of the sorted eigenvectors and eigenvalues provides additional flexibility in selecting the target dimensionality of the latent space.

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

$$\{\mathbf{r}_t\}, \quad t \in [0, dt] \quad (9)$$

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.

² $\delta_{ij} = 1$ if $i = j$, 0 otherwise.

3.3 Smoothing

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 some cases be eliminated entirely.

4. RESULTS

The two algorithms were applied to a variety of motion datasets. Our experiments are divided into three sections. The first section deals with the analysis of walking motion, which exhibit clear periodic properties. The goal in this case is the identification of the various segments that constitute the periodic motion. 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. The third section of our results deals with the more challenging case of motion recovery under noisy or imperfect sensing.

In the first two sections of our results, we present a comparison between the basic model-based algorithm, and the model-free, latent space algorithm, in the context of motion segmentation and activity detection. We show that the latent space algorithm can recover points and regions of interest that closely match the corresponding points of the model-based algorithm. In the third part of this Section, we analyse periodic motions that were captured in a noisy hospital environment, which impacted the performance of our sensors. We demonstrate that the latent space algorithm can recover the qualitative structure of these motions even when some of the sensor readings are omitted.

4.0.1 Ground truth data. Given that we evaluate our algorithms on real, unprocessed motion capture data, the most effective way to validate the correctness of our results is through visual comparison. We have uploaded videos for the datasets presented in this section on the supporting material page (<http://www.specknet.org/publications>). For each dataset, we provide two videos: the raw video of the motion sequence as captured by the sensors, and a second visualisation with the segmentation points and activity intervals overlaid.

4.1 Evaluating periodic motions

For this set of experiments, the subject was assigned with the set task of walking a few steps forward, turning around 180° , and walking 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. Figure 2 illustrates the difference between motion on the transverse and the sagittal plane, for the left and right knee joints. In the second subfigure, 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, even though the unit is the same (centimeters) on

both planes, the curves in each graph represent motion along different axes.

The *model-based* segmentation points 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.

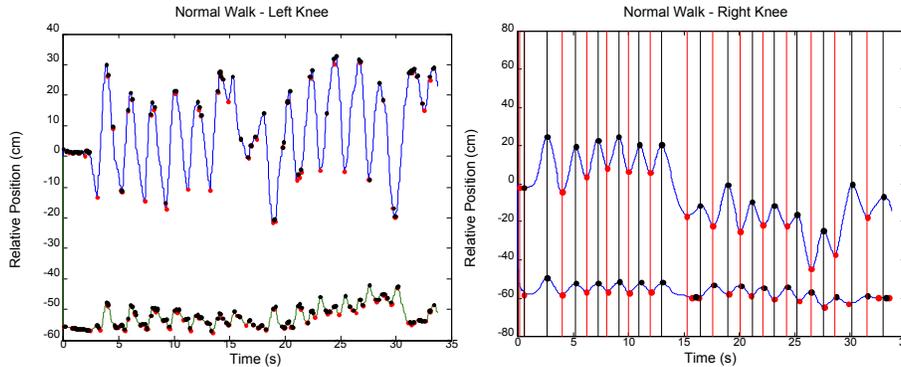


Fig. 2. Left: Left knee motion on transverse (top curve) and sagittal (bottom curve) plane (smoothing factor = 5). Right: Overlaid segmentations for right knee motion, transverse (top curve) and sagittal (bottom curve) plane (smoothing factor = 40)

In the first half of the sequence, the minima and maxima on the two planes are closely aligned. The 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, so 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-right is an example of this failure. Model-based segmentation correctly identifies all the extrema 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 Figure 2 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, are arbitrary and

must be made by the human designer. It is also sometimes necessary to perform different types of comparison in order to verify segmentation results. For example, in Figure 2-right, 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.

Latent space segmentation avoids these problems by providing an automated and principled way of characterising periodic motion. The method compactly 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 Figure 2, which suffered from the computation of several false positives. We first use the Isomap procedure 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.

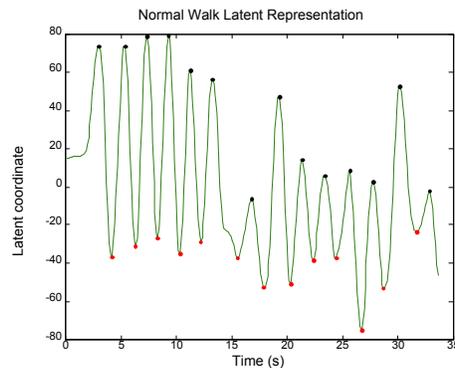


Fig. 3. Latent space segmentation of the normal walking motion analysed in Figure 2

Figure 3 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 *units-free*, 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 all observed spatial dimensions represents coordinates in centimeters. 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 quantitative 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.

As Figure 4 demonstrates, the latent space algorithm manages to avoid false positives. Furthermore, even though it was previously stated that the latent coordinate is units-free, it is still possible to make qualitative observations about the nature of the motion. This is

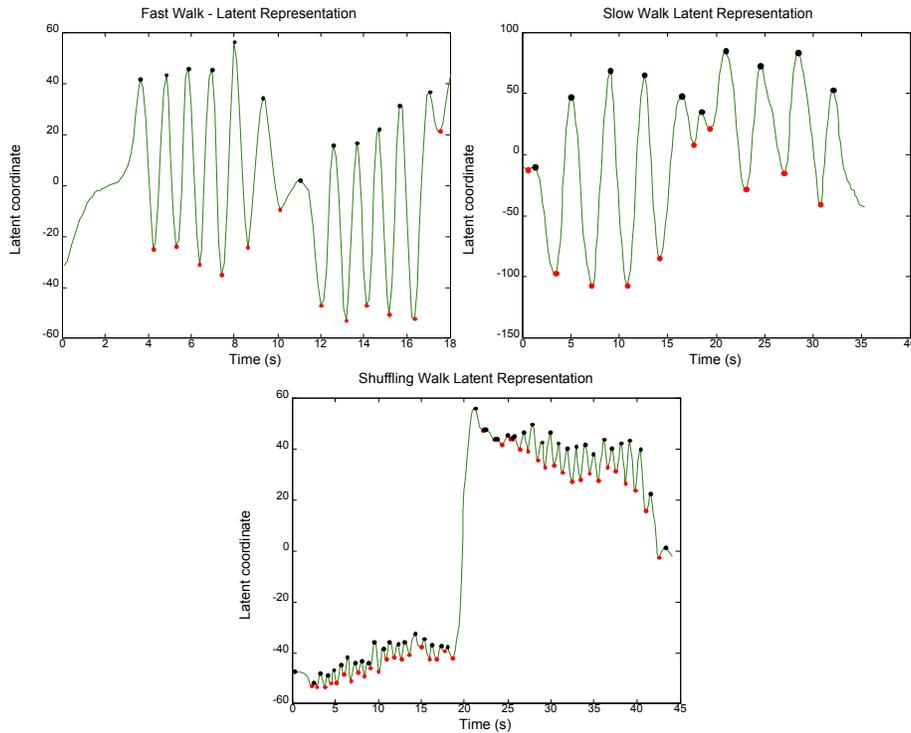


Fig. 4. Fast walk (top left), slow walk (top right) and shuffling walk (bottom) - latent space segmentation

most apparent in Figure 4-bottom, 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 Algorithmic Extensions. The examples discussed so far consider a one-dimensional latent space representation of a walking sequence. This target dimensionality was chosen because it simplifies the segmentation procedure; all that is required is to look for local minima and maxima, which is trivial in a single dimension. However, the latent space algorithm can be extended to higher dimensions, by replacing the min-max procedure with detection of *critical points* in the curve of the motion. This approach may also lead to a more specific characterisation of the segmentation points, for example, by considering the sign of higher-order derivatives.

Furthermore, Isomap may be replaced by other dimensionality reduction algorithms to yield similar qualitative results. As a comparative method, we have chosen the Gaussian Process Dynamical Model algorithm (GPDM) [Wang et al. 2006], a probabilistic algorithm designed for temporally extended sequences. Figure 5 shows the two-dimensional representation of the fast walk of Figure 4, for both Isomap and GPDM. Both visualisations are comparable to the one-dimensional representation of the motion, although the GPDM takes significantly longer to generate this result (around 48 seconds, compared to 3 for

Isomap). The two-dimensional plot also highlights some of the qualitative features of the motion, such as the reversal in the walking direction halfway through the sequence, which is reflected by the transition of the curve to a different plane.

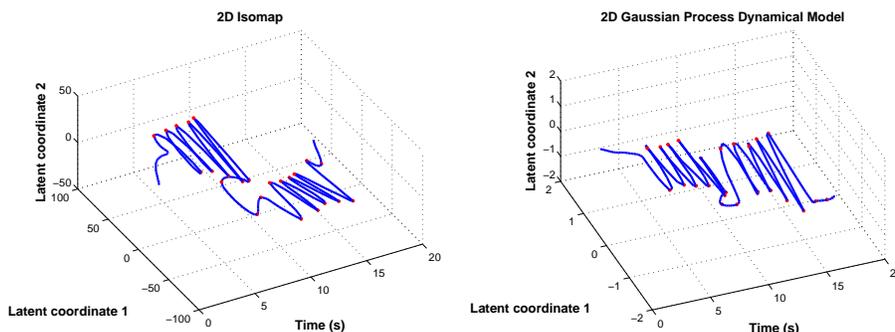


Fig. 5. Two dimensional representation for fast walk of Figure 4 - Isomap (left) and Gaussian Process Dynamical Model (right)

The identification of critical points in the latent space of the motion is an illustration of how complex sequences may be simplified in lower dimensions. These manifolds provide simplified representations that can be analysed without extensive parameter tuning. The identification of local minima and maxima for walking period detection can be lifted to other motion tasks, whose topology may be more intricate and requiring more sophisticated critical point interpretation. However, as Figure 5 suggests, the underlying qualitative features and dynamics may be maintained and recovered, regardless of the precise target dimensionality and manifold learning technique that is being used.

4.1.5 Walking period and error estimation. As seen on Figure 3, the latent space 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 Figure 2 that over-segment the motion.

Nevertheless, it is important to represent any errors between the two methods in numerical terms. Figure 6(a) presents the error bars for the latent segmentation points computed in Figures 3-4. 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:

$$\mathbf{err}(lmax) = \min_{m \in SPmaxs} |lmax - m| \quad (10)$$

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

$$\text{err}(lmin) = \min_{m \in SP_{mins}} |lmin - m| \quad (11)$$

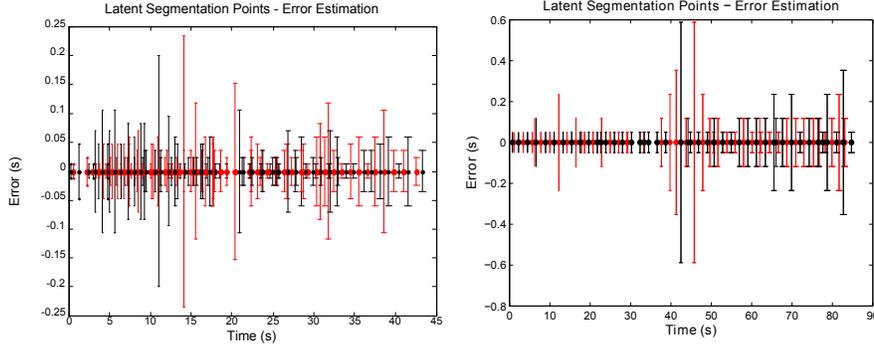


Fig. 6. *Left*: Similarity error estimation for the computed latent segmentation points of Figures 3-4 (red-minima, black-maxima). *Right*: Error estimation for a further three walks performed by a second subject (mean similarity error = 0.0462s). Note that, although the unit in both axes is seconds, the scales of the axes are *not* equal.

Table I summarises the walking period and 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, and the corresponding similarity error percentage is given. The computed periods accurately reflect the differences between the different walking 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 (<http://www.specknet.org/publications>) provide a means of visual validation and verification of these results.

For validation purposes, the error bars for three walks performed by a second healthy male subject are given in Figure 6(b). The latent space segmentation for these walks is given in Figure 21 (Appendix A).

	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 I. 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 sequence has been 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.

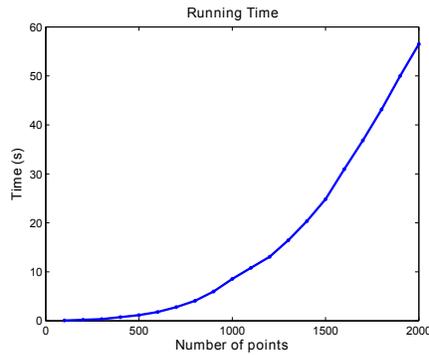


Fig. 7. Running time of the latent space segmentation algorithm. Experiments were run on Dell Optiplex 780 Desktop machine running Windows Vista Home Basic, with an Intel Core 2 Duo Processor and 4GB RAM.

4.1.6 Running time evaluation. 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. As explained in Section 3.2.1, the Isomap algorithm estimates this manifold by first constructing a neighbourhood graph over the high-dimensional points. This means that each data point needs to be compared against all others, which unfortunately scales exponentially with the size of the dataset. By contrast, the time taken to process the low-dimensional sequence is linear to the number of points, so negligible compared to the learning step. Figure 7 visualises the running times for datasets of various sizes. When the sets do not exceed 500 points, a one-dimensional manifold can be learned in less than half a second, keeping the overall complexity of the algorithm low. However, if this limit is exceeded, the process may take up to a few seconds or even minutes, for large datasets close to or greater than 2,000 points.

A simple yet effective solution is to *prune* our datasets. Since the Orient sensors provide orientation data at a frequency exceeding 100Hz, more data is available than was necessary to compute segmentation points. As such, it is possible to prune the datasets by a factor of 10 or even 20, without losing much crucial information. In the context of dimensionality reduction, pruning is closely related to the family of landmark-based techniques (such as Landmark Isomap), which consider only a subset of a given dataset in order to compute an embedding.

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 is the identification of the intervals where walking occurs, and correspondingly 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. Data was captured in two distinct phases. During the first phase, two patients (labeled 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 identify points that match a predefined movement intensity threshold. This threshold has been defined in order to detect possible regions of motion activity.

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.

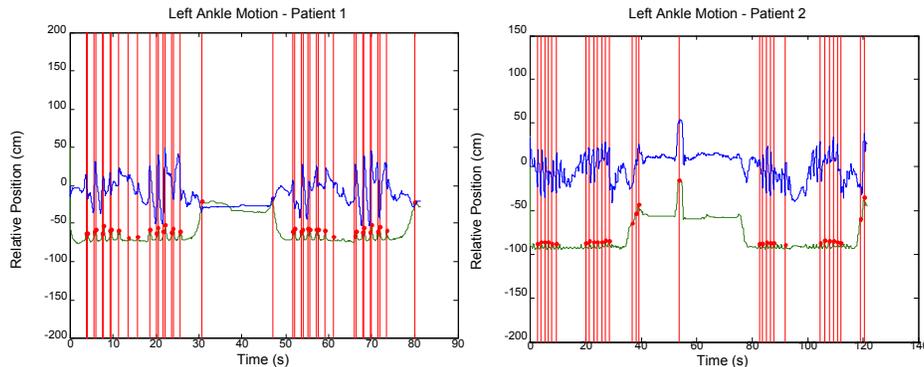


Fig. 8. Left ankle motion for Patient 1 (left) and Patient 2 (right) - transverse (top blue curve) and sagittal plane (bottom green curve)

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 8-right. 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 8-right 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.

4.2.3 Latent space activity detection. When evaluating the motion of elderly people, 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.

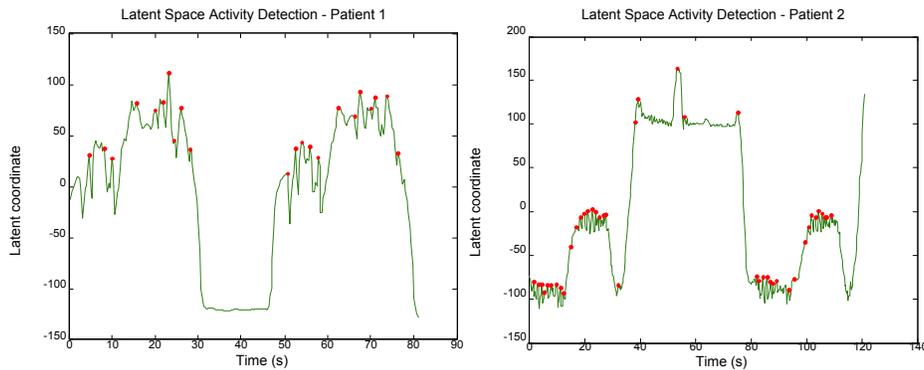


Fig. 9. Latent space activity detection - Patient 1 (left) and Patient 2 (right)

Figure 9 demonstrates 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 8-right. Moreover, it is now possible to define four distinct clusters of intense motion in Figure 9-right, which were not as clear in Figure 8.

4.2.4 Validation. The activity detection algorithms were also applied to data from a third patient, who was diagnosed with severe mobility problems.

Figure 10(a) indicates a different behaviour to the one identified in Figure 9. Because of the patient’s mobility problems, it is now more difficult to discern intervals of activity. This feature is reflected in Figure 10(a). 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 10(b) shows the latent space representation of the above sequence, which marks more clearly the region where motion occurs in Figure 10(a).

4.3 Qualitative analysis with fewer sensor nodes

Wireless inertial motion capture is an inherently complex procedure, involving the interaction and fusion of several sensor readings. Moreover, the nature of the utilised sensors

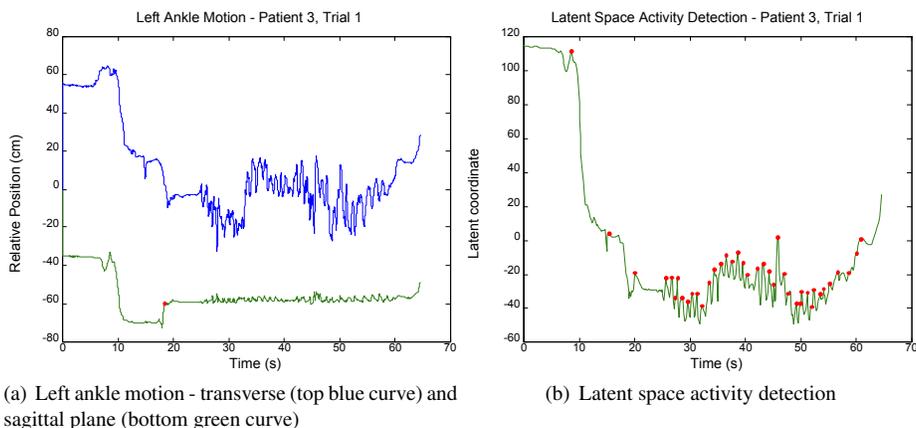


Fig. 10. Patient 3 motion analysis

may also be diverse; in our experiments, our devices comprise triaxial gyroscopes, magnetometers, and accelerometers. The performance of these sensors is often strongly correlated with the environment where the motion is captured. For example, magnetometers are very sensitive to iron and steel, yielding inconsistent or imprecise measurements in the presence of these materials. This degeneracy impacts the overall quality of the captured motion, making its description and analysis difficult.

It is therefore essential for motion analysis techniques to be robust with respect to imperfect or incomplete raw sensory data. Furthermore, an algorithm that can cope with fewer sensor nodes can reduce the energy consumption of a wireless sensor network significantly.

In Sections 4.1 and 4.2, we demonstrated how our latent space algorithm can summarise large motion sequences effectively, without the need of excessive preprocessing such as smoothing. In this section, we show how this approach can deal with problematic sensor devices. In particular, we analyse sequences where some of the devices were suffering from *constant drift* errors, leading to dropped packets that severely impacted the quality of the captured motion.

4.3.1 Normal walk segmentation revisited. We first revisit the normal walk example of Figure 3, where 4 devices had been used in total, leading to a 12-dimensional motion vector. In this section, we progressively reduce this number, and determine its effect on the performance of the latent space segmentation algorithm.

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 3. However, even with fewer sensor nodes, the latent space algorithm avoids the over-segmentations produced by the model-based variant of Figure 2, which also uses data from one sensor only.

4.3.2 Experiments with hospital data. We now consider a more practical example, where the removal of motion data is dictated by the poor quality of sensor readings. The datasets were captured in the Astley-Ainslie Hospital Mobile Gait Laboratory in Edinburgh. The subjects were young children suffering from cerebral palsy or similar disorders that impact their gait. As in previous sections, the Orient devices were used to capture the

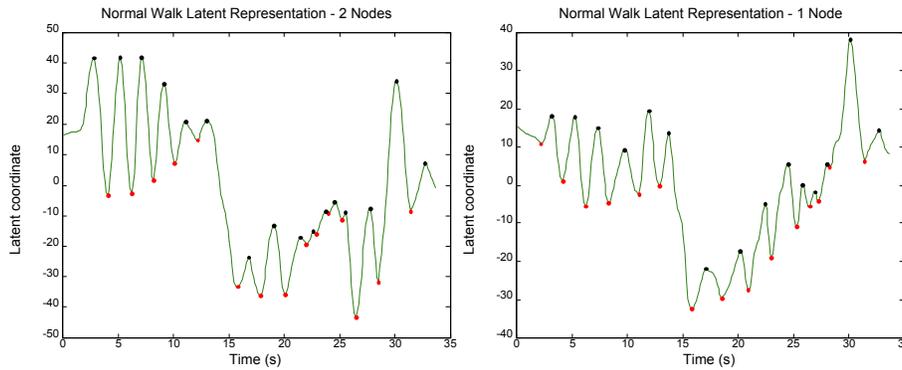


Fig. 11. Left: Latent space representation of a normal walk using 2 sensor nodes (right/left knee). Right: Latent space representation of the same walk using 1 sensor node (right knee)

data; however, because of the nature of the experimental environment (high concentration of steel and iron), the magnetometers were switched off. Unfortunately, these features also impacted the gyroscopes and the accelerometers, whose output was noisy and inconsistent.

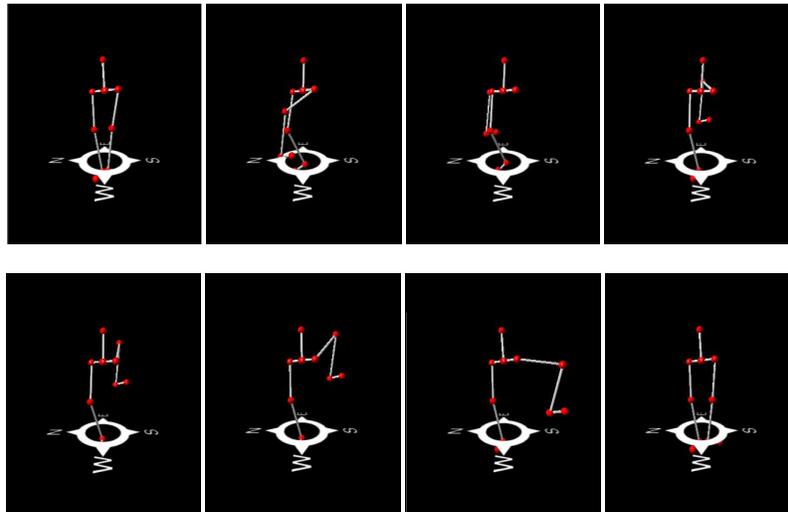


Fig. 12. Snapshots from a single step of Noisy Motion Sequence 1, illustrating the effect of the noisy data.

In the first dataset (Noisy Motion Sequence 1), the subject walked about around a small area, with a total of 6 devices placed on his legs (left/right thigh, knee and ankle). The effect of noisy sensor data is strongly apparent in Figure 12, which is a visualisation of a single step taken by the subject. Most snapshots display physically implausible poses for the left leg, with limbs crossing over each other or forming unnatural angles between them.

To quantify this irregularity, Figure 13 plots the left and right thigh z-axis rotational rates, in radians/s, as measured by the corresponding gyroscopes. This figure partly explains the irregular motion observed on the left leg during the course of the motion se-

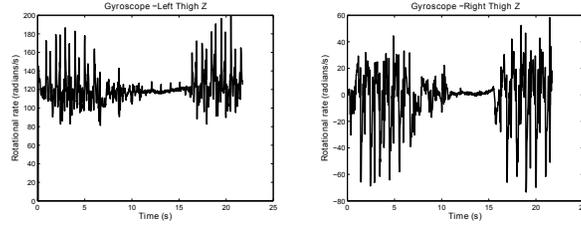


Fig. 13. Left thigh (faulty) and right thigh (normal) z-axis rotational rate for Noisy Motion Sequence 1

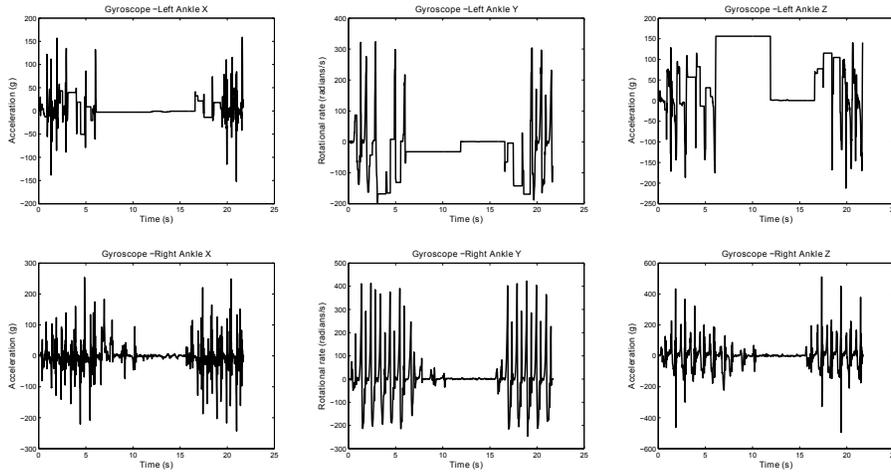


Fig. 14. Noisy Motion Sequence 1 - rotational rates for nkle joints. *Top*: Left ankle (faulty). *Bottom*: Right ankle (stable).

quence - while the right thigh sensor yields consistent values that are centred around 0 radians/s, the left thigh sensor suffers from a constant drift that shifts these values in excess of 100 radians/s.

Furthermore, the top halves of Figures 14-15 depict a different type of error, as observed in the rotational rates and accelerations measured by the two ankle devices. Instead of suffering from a constant drift like the thigh gyroscopes, the sensors drop several packets over the course of the capture. The flat regions for the left ankle (Figures 14-top and 15-top) indicate the intervals where these packets were dropped; by contrast the right ankle sensors (Figures 14-bottom and 15-bottom) produce consistent measurements. The existence of a similar error in both the gyroscopes and the accelerometer indicates an overall problem with the device placed on the left ankle.

To evaluate the performance of our algorithm under these harder constraints, we computed the latent space representation of Noisy Motion Sequence 1. Figure 16 - left shows the representation of the sequence when all 18 dimensions of the motion are used. In Figure 16 - right, we have removed the children of the problematic nodes, namely the left knee and left ankle motion data. This Figure is more consistent with the stable readings of Figures 14 and 15, which include two clear clusters at the start and the end, and little

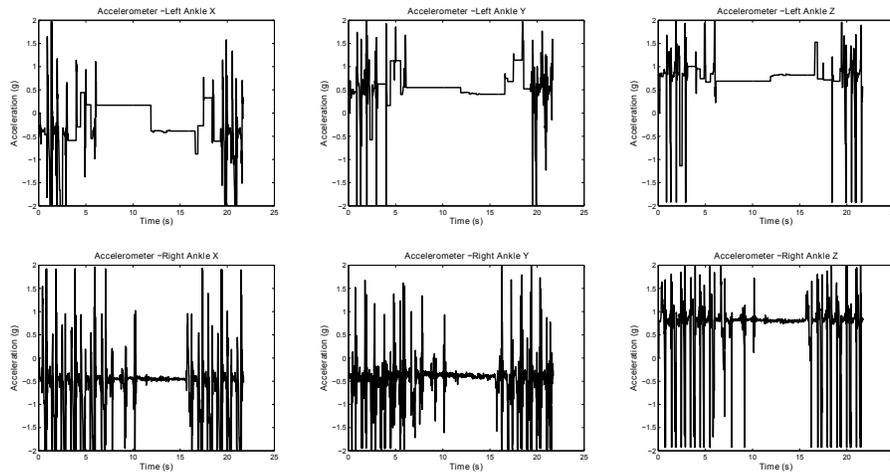


Fig. 15. Noisy Motion Sequence 1 - measured accelerations for the ankle joints. *Top:* Left ankle (faulty). *Bottom:* Right ankle (stable).

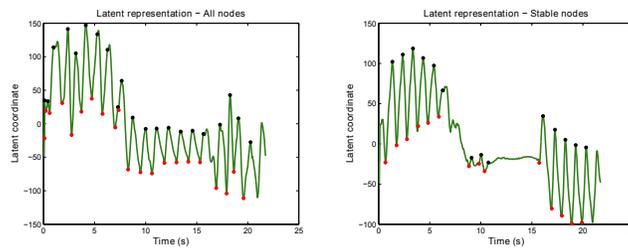


Fig. 16. Latent space representation for Noisy Motion Sequence 1 - all sensors (left) and stable sensors only (right)

motion in the middle. By contrast, Figure 16 - left is impacted by the constant drift of the left thigh gyroscope, thus representing an ongoing motion activity throughout the capture.

Figure 17 plots the latent representations of another motion sequence from the same subject. The gyroscope and accelerometer readings for all 18 axes are also plotted for verification. The stable motion vector produces the results most closely matching the average gyroscope and accelerometer readings, while also yielding a smoother motion representation. Moreover, the vector consisting of all sensor readings identifies a periodic motion during the first half of the capture, in contrast to the raw sensor readings.

4.3.3 Validation. Finally, we validate the results of this section by plotting results from captures involving a different patient (Noisy Motion Sequences 3 and 4). The captures were held in the same location as the previous ones, but on a different date and with different sensor devices. In this particular case, the problematic devices were identified as being located on the ankles of the patient. Figure 19 illustrates this effect, by visualising the gradual drift in both ankle positions over the course of the capture. As in Figures 16-17, Figures 18 and 20 show the full and stable latent space representations, as well as the raw

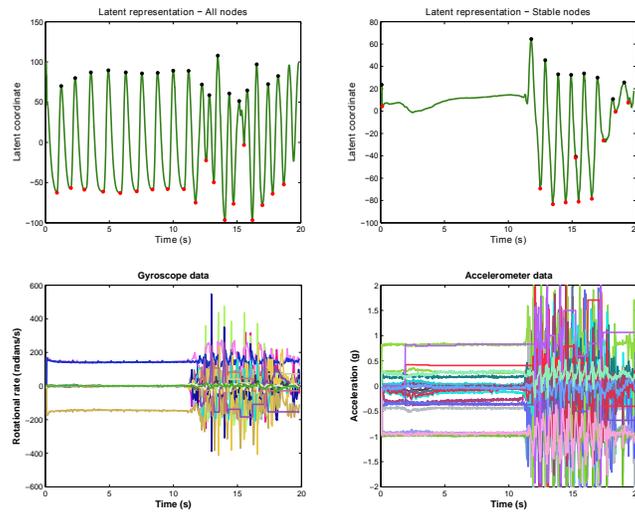


Fig. 17. Latent space representation for Noisy Motion Sequence 2 - all sensors (top left) and stable sensors only (top right), all gyroscope readings (bottom left) and all accelerometer readings (bottom right)

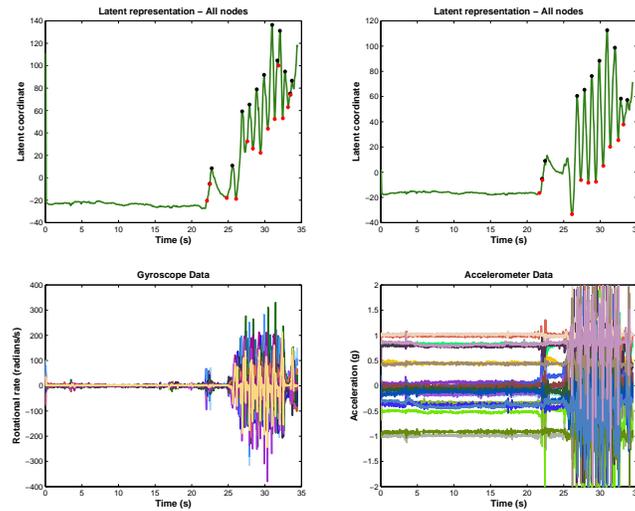


Fig. 18. Latent space representation for Noisy Motion Sequence 3 - all sensors (top left) and stable sensors only (top right), all gyroscope readings (bottom left) and all accelerometer readings (bottom right)

gyroscope and accelerometer readings.

Compared to the previous set of experiments, there is a less marked difference between the full and stable latent space representations. This discrepancy is due to two factors. First, the nature of the sensory error is closer to the one displayed in Figure 13 than to the one in Figures 14 and 15; that is, there is a consistent scaling error in the readings and not a loss of packets. Second, the error is located solely on the ankle devices and not on the thigh

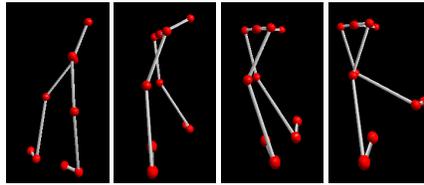


Fig. 19. Snapshots from consecutive steps in Noisy Motion Sequence 4, illustrating the accumulating error in the estimation of the ankle postures.

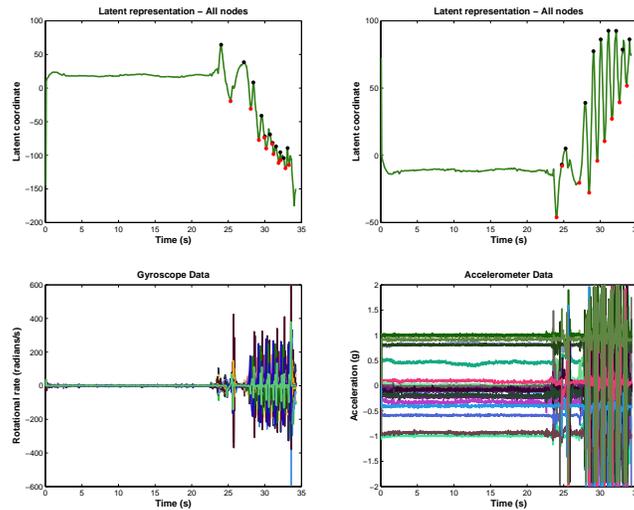


Fig. 20. Latent space representation for Noisy Motion Sequence 4 - all sensors (top left) and stable sensors only (top right), all gyroscope readings (bottom left) and all accelerometer readings (bottom right)

device as well, as was the case in Noisy Motion Sequences 1-2. As the ankles lie near the bottom of the tree-like body model structure, they impact the relative position estimates of fewer joints than the thighs. Thus, the overall latent representation of the periodic motion does not degrade significantly, even when problematic devices are considered.

5. CONCLUSION

5.1 Strengths

An unsupervised, latent-space algorithm has been presented for analysing a person’s gait using the on-body wireless Orient motion capture system. The algorithm was compared to a model-based variant for the purposes of period identification in periodic walks. Moreover, the latent space algorithm was tested 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 latent space algorithm was shown to be robust with respect to noisy or incomplete sensor readings. This is a commonly encountered constraint in wireless sensor network applications, as many experimental environments are inherently unsuited to inertial sensing. Furthermore, the Orient devices we use in our experiments operate in a *distributed* manner, so joint positions are computed by combining the readings of multiple sensors. Thus, it is important for motion analysis algorithms to be insensitive to drifts or recurring measurement errors that may arise in some of the deployed sensors. The results demonstrate that the latent space representation can recover the *qualitative* structure of the motion sequence, in the face of varying sensor error types (constant drift, dropped packets) and noisy sensor locations (thigh, ankle). This indicates that the algorithm can be used as a smoothing tool, through which the impact of faulty devices on the overall motion quality can be minimised.

5.2 Limitations and future work

The computational analysis of our algorithm showed that the running time scales exponentially with the size of the analysed dataset. 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. However, even with dataset pruning (Section 7), it is still not possible to apply the latent space segmentation online. The reason is that the algorithmic procedure of manifold learning requires that *neighbourhood 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 learned in a supervised manner and used to embed points to the latent space online. Nonetheless, learning mappings that generalise well to unseen data points is a difficult task, and one that few existing manifold learning techniques address adequately.

When evaluating the latent space algorithm on noisy motion sequences, the elimination of the problematic devices was done manually, by first visually inspecting all the sensor readings. It would be more beneficial if this procedure was automated, while being more closely integrated with the low-dimensional aspect of the procedure. For example, rather than always keeping the largest eigenvector during the embedding, it would be worth considering other large eigenvectors as candidates for better summarising the motion. If such an association between sensor data quality and spectral analysis of the motion can be formalised, it might even be possible to apply the latent space algorithm without removing any noisy data beforehand.

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. Alternatively, the target dimensionality of the latent space could be adjusted to reflect the multiple modes of a motion sequence. In such a case,

we would also need to go beyond simple min-max heuristic, and exploit the correlation between the different dimensions of the resulting embedding.

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 categorise motions. Nevertheless, the main motivation behind our proposed approach is that good quality motion segmentation is a prerequisite for 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). The datasets analysed in Section 4.3 were captured by Speckled Computing research fellow Smita Sasindran.

REFERENCES

- ALBU, A. B., BERGEVIN, R., AND QUIRION, S. 2008. Generic temporal segmentation of cyclic human motion. *Pattern Recogn.* 41, 1, 6–21.
- BISHOP, C. M. 2007. *Pattern Recognition and Machine Learning (Information Science and Statistics)*, 1st ed. 2006. Corr. 2nd printing ed. Springer.
- BLAKE, A., ISARD, M., AND REYNARD, D. 1995. Learning to track the visual motion of contours. *Artificial Intelligence* 78, 101–134.
- CHEN, J., KWONG, K., CHANG, D., LUK, J., AND BAJCSY, R. 2005. Wearable sensors for reliable fall detection. In *Proceedings of the 2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*.
- DE SILVA, V. AND TENENBAUM, J. B. 2003. Global Versus Local Methods in Nonlinear Dimensionality Reduction. In *Advances in Neural Information Processing Systems 15*. Vol. 15. 705–712.
- GUENTERBERG, E., GHASEMZADEH, H., AND JAFARI, R. 2009. 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*. IEEE Computer Society, Washington, DC, USA, 339–344.
- GUENTERBERG, E., GHASEMZADEH, H., LOSEU, V., AND JAFARI, R. 2009. 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*. Springer-Verlag, Berlin, Heidelberg, 145–158.
- GUENTERBERG, E., OSTADABBAS, S., GHASEMZADEH, H., AND JAFARI, R. 2009. 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*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), ICST, Brussels, Belgium, Belgium, 1–7.
- JEAN, F., BERGEVIN, R., AND BRANZAN-ALBU, A. 2005. Body tracking in human walk from monocular video sequences. In *Proceedings of the Second Canadian Conference on Computer and Robot Vision (CRV)*. Victoria, BC, Canada, 144–151.
- LAPTEV, I., BELONGIE, S. J., PREZ, P., AND WILLS, J. 2005. Periodic motion detection and segmentation via approximate sequence alignment. In *In Proceedings of the International Conference on Computer Vision*.
- LAWRENCE, N. 2003. Gaussian process latent variable models for visualisation of high dimensional data. In *NIPS*.
- LEI, Y.-K., XU, Y., ZHANG, S.-W., WANG, S.-L., AND DING, Z.-G. 2010. Fast isomap based on minimum set coverage. In *Advanced Intelligent Computing Theories and Applications. With Aspects of Artificial Intelligence*. Lecture Notes in Computer Science, vol. 6216. Springer Berlin / Heidelberg, 173–179.

- LV, F. AND NEVATIA, R. 2006. Recognition and segmentation of 3-d human action using hmm and multi-class adaboost. In *ECCV 2006*, A. Leonardis, H. Bischof, and A. Pinz, Eds. Lecture Notes in Computer Science, vol. 3954. Springer Berlin / Heidelberg, 359–372.
- MACDORMAN, K. F., CHALODHORN, R., AND ASADA, M. 2004. Periodic nonlinear principal component neural networks for humanoid motion segmentation, generalization, and generation. *Pattern Recognition, International Conference on 4*, 537–540.
- MIYAZAKI, S. 1997. Long-term unrestrained measurement of stride length and walking velocity utilizing a piezoelectric gyroscope. In *IEEE Transactions on Biomedical Engineering*. Vol. 44. 753–759.
- QUIRION, S., BRANZAN-ALBU, A., AND BERGEVIN, R. 2005. Skeleton-based temporal segmentation of human activities. In *Proceedings of the 13th International Conference in Central Europe on Computer Graphics (WSCG 05)*. 145–148.
- RABINER, L. R. 1989. A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE*. 257–286.
- ROWEIS, S. T. AND SAUL, L. K. 2000. Nonlinear dimensionality reduction by locally linear embedding. *Science* 290, 5500 (December), 2323–2326.
- RUI, Y. AND AN, P. 2000. Segmenting visual actions based on spatio-temporal motion patterns. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- SABATINI, A., MARTELLONI, C., SCAPELLATO, S., AND CAVALLO, F. 2005. Assessment of walking features from foot inertial sensing. In *IEEE Transactions on Biomedical Engineering*. Vol. 52. 486–494.
- TENENBAUM, J. B., DE SILVA, V., AND LANGFORD, J. C. 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 5500 (Dec), 2319–2323.
- THANGALI, A. AND SCLAROFF, S. 2005. 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.
- VALTANOS, A., ARVIND, D. K., AND RAMAMOORTHY, S. 2010. Comparative study of segmentation of periodic motion data for mobile gait analysis. In *Wireless Health 2010. WH '10*. ACM, New York, NY, USA, 145–154.
- VAN DER MAATEN, L., POSTMA, E., AND VAN DEN HERIK, H. 2009. Dimensionality reduction: A comparative review. Tech. Rep. TiCC-TR 2009-005, Tilburg University.
- WANG, J. M., FLEET, D. J., AND HERTZMANN, A. 2006. Gaussian process dynamical models. In *NIPS*. Vol. 18. 1441–1448.
- YANG, A. Y., IYENGAR, S., SASTRY, S., BAJCSY, R., KURYLOSKI, P., AND JAFARI, R. 2008. Distributed segmentation and classification of human actions using a wearable sensor network. In *in Proceedings of the CVPR Workshop on Human Communicative Behavior Analysis*.
- YOUNG, A. D., LING, M. J., AND ARVIND, D. K. 2007. Orient-2: a realtime wireless posture tracking system using local orientation estimation. In *Proceedings of the 4th workshop on Embedded networked sensors. EmNets '07*. ACM, New York, NY, USA, 53–57.

A. VALIDATION DATASETS FOR SECOND HEALTHY SUBJECT

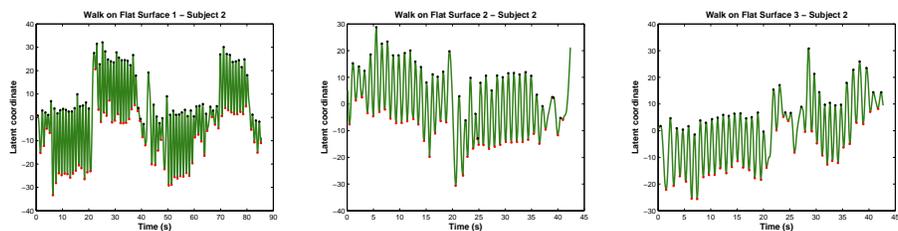


Fig. 21. Latent space segmentation for three long walking sequences performed by a second healthy subject