Posture Recognition and Postural Transition Detection Using Bed-Based Pressure Sensor Arrays

by

Nicholas Foubert, B. Sc.

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of

Master of Applied Science in Biomedical Engineering

Ottawa-Carleton Institute for Biomedical Engineering (OCIBME)
Department of Systems and Computer Engineering

Carleton University
Ottawa, Ontario, Canada, K1S 5B6

April 2010

Your file Votre référence
ISBN: 978-0-494-68652-2
Our file Notre référence
ISBN: 978-0-494-68652-2

The undersigned recommend to the Faculty of Graduate Studies and Research acceptance of the thesis

Posture Recognition and Postural Transition Detection Using Bed-Based Pressure Sensor Arrays

submitted by

Nicholas Foubert, B. Sc.

in partial fulfillment of the requirements for the degree of

Master of Applied Science in Biomedical Engineering

Chair, Howard Schwartz, Department of Systems and Computer Engineering

---

Thesis Supervisor, Rafik Goubran

Thesis Supervisor, Aysegul Cuhadar

Carleton University
April 2010

Abstract

This thesis focuses on the problem of using bed-based optical pressure sensor arrays to unobtrusively recognize if a person is sitting, lying in a particular posture, or transitioning from lying to sitting while in bed. Recognizing postures and postural transitions enables meaningful segmentation of pressure signals as a preprocessing step before applying posture-specific algorithms. Detecting lying postures is also relevant for prevention and treatment of ailments including obstructive sleep apnea, sudden infant death syndrome, and gastro-esophageal reflux. In addition, postural transitions such as sit-to-stand and lie-to-sit are commonly by used by health care professionals to assess patient mobility. We show that, using pressure data from a bed-based optical pressure sensor array, sitting and lying postures can be reliably recognized and that lie-to-sit postural transitions can be detected with low miss rate and accuracy commensurate with video analysis.

Acknowledgements

First and foremost I would like to thank Dr. Rafik Goubran. Dr. Goubran has provided me immeasurable support throughout my masters program. If all professors were like Dr. Goubran, everyone would aspire to enter graduate school! Dr. Aysegul Cuhadar, my co-supervisor, provided much advice and constructive criticism of my work. Although not nominally one of my supervisors, Dr. Frank Knoefel provided much needed guidance and great opportunities to engage with students in the medical field. I had the pleasure of contributing some of my work to two conference submissions prepared by Dr. Knoefel and his students Pam Carlson and Shawn Mondoux.

I would also like to thank my parents, Jeannette and Dwight, and my brother Patrick for all of the caring support that I've benefited from over the years. My partner Kellie probably deserves a medal for putting up with me over the past two years. I'm very thankful for all of her support too; she took good care to keep me from going off the deep end on more than a few occasions.

Thanks also to all of my fellow students in the Digitial Signal Processing lab at Carleton and particularly those working with the TAFETA group.

Thanks to the Ontario Government and Carleton University for supporting me financially through scholarships including the Ontario Graduate Scholarship (OGS). This research has also been funded by the Natural Sciences and Engineering Research Council (NSERC).

Contents

List of Figures

List of Tables

List of Abbreviations and Symbols

Abbreviations

Abbreviation	Definition
DWT	Discrete wavelet transform
F	Fetal posture
FN	False negative
FP	False positive
FSR	Force sensitive resistor
kNN	k-nearest neighbour
L	Lateral (decubitus) posture
LatCP	Lateral center of pressure
LatV	Lateral variance
LF	Left fetal posture
LiSi	Lie-to-sit postural transition
LiSt	Lie-to-stand postural transition
LL	Left lateral (decubitus) posture
LonCP	Longitudinal center of pressure
LonV	Longitudinal variance
MAS	Motor assessment scale
MSAS	Mobility scale for acute stroke patients

Abbreviation	**Definition**
NAS	Number of active sensors
NN	Nearest neighbour
Pr	Prone posture
RBF	Radial basis function
RF	Right fetal posture
RL	Right lateral (decubitus) posture
SD	Subject-dependent
SI	Subject-independent
SIDS	Sudden infant death syndrome
Si	Sitting up
SiLi	Sit-to-lie postural transition
SiSt	Sit-to-stand postural transition
SSV	Sum of sensor values
StLi	Stand-to-lie postural transition
StSi	Stand-to-sit postural transition
Su	Supine posture
SVM	Support vector machine
TAFETA	Technology Assisted Friendly Environments for the Third Age
TN	True negative
TP	True positive
WNAS	Weighted number of active sensors
WSSV	Weighted sum of sensor values

Symbols

Symbol	Definition
$\mathbf{w}$	Normal vector of a separating hyperplane
b	SVM bias
$\phi(\cdot)$	Kernel function
ξ_i	SVM error tolerance
C	SVM misclassification penalty parameter
γ	SVM hyperparameter when using an RBF kernel
τ	Instance class
t	Discrete time (sample) index
$S_i[t]$	State of ith sensor at time t
x_i	longitudinal coordinate of ith sensor
y_i	lateral coordinate of ith sensor
$s_i[t]$	Value of the ith sensor at time t
N	Total number of sensors
A	Step function
w	Sensor weighting function
L	Length of pressure sensitive mat
λ	Sensor weight
$\Delta M[t]$	motion vector sequence
T_M	threshold on motion vector sequence
L_W	Window length for averaging low-pass filters
σ	sequence of movement segments
$\Phi[t]$	static activity (sitting or lying) at time t

Chapter 1

Introduction

The Canadian Special Senate Committee on Aging released a report in April of 2009 expressing a need to understand how new technologies may help mitigate the effects of limited health care resources in a time when birth rates are decreasing, life spans are increasing, and the overall population is aging [1]. In this context, automated and unobtrusive monitoring of health status and activities of daily living in elderly patients is becoming an increasingly important area of study as health care resources are spread thinner to support a growing senior population [2].

Smart home environments are actively being studied to test the capabilities of various sensing modalities to provide accurate and salient information about patients both at home and in clinical environments. One particular subject area which is receiving increasing attention in the research community is pressure sensitive home environments [3]. This thesis belongs squarely within that area of research, investigating the feasibility of using pressure sensor arrays placed beneath a bed mattress to recognize a person's physical context while they are in bed. It focuses specifically on two related activities: recognizing a person's posture and detecting their postural transitions using a bed-based pressure sensor array.

Figure 1.1 depicts the basic flow of information from the person on the bed to

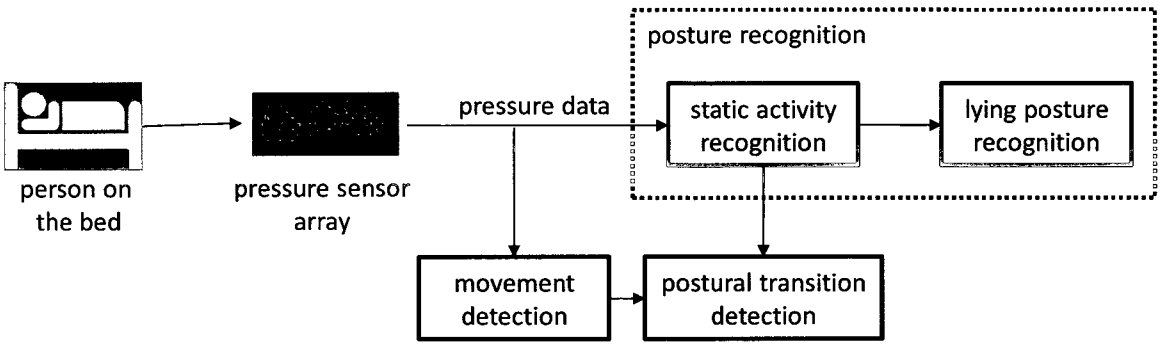

Figure 1.1: Bed-based pressure sensing system for recognizing postures and postural transitions.

posture recognition and postural transition detection. There are three main parts to the system that are studied in this work:

1. **static activity recognition** determines whether the person is sitting or lying on the bed;

2. **lying posture recognition** determines the lying posture (e.g. lying on their back) assumed when the person is lying; and

3. **postural transition detection**, in the context of bed-based pressure sensor arrays, detects when the person on the bed transitions from sitting to lying or lying to sitting.

Static activity and lying posture recognition, which in this work are collectively referred to as **posture recognition**, are investigated in chapter 4. The postural transition detection system takes input from both the static activity recognizer and a simple movement detector. Movement detection and postural transition detection are discussed in chapter 5. Note that, for the sake of brevity, the static activity recognition system, lying posture recognition system, and postural transition detection system are also referred to as the **static activity recognizer**, **lying posture recognizer**, and **postural transition detector**, respectively.

1.1 Statement of the problem

This thesis aims to answer the following research question: is it feasible to use a pressure sensor array, placed beneath a mattress typical of Canadian hospitals, to automatically, unobtrusively, and reliably 1) recognize a person's static activity (i.e. sitting and lying) on the bed, 2) recognize a person's lying posture when lying, and 3) both recognize and record the duration of the person's lie-to-sit postural transitions. In this context, the system is automatic if, after an initial period of system training, it performs the recognition tasks without any input from a system operator. The system is unobtrusive if it does not require the attachment of any device to the person's body. Reliability is assessed by comparison to external standards of reliability, as defined in the subsequent chapters.

The work presented herein differs from the related work discussed in chapter 2 in the following ways:

1. previous studies using pressure sensitive mats for posture classification, as presented in section 2.3.2, employ custom built arrays of force sensitive resistors (FSRs), while the present work uses an optical pressure sensing technology;

2. this study uses a mattress typical of Canadian hospitals, which is thicker than the futon-type mats used in related work discussed in section 2.3.2 and specifically designed to more effectively distribute pressure for comfort;

3. a more complete set of lying postures is considered, including particularly the left/right fetal posture, compared with those mentioned in section 2.3.2; and

4. the methods of lie-to-sit postural transition detection reviewed in section 2.3.3 all use wearable sensing technologies, so to the best of our knowledge this is the first work investigating lie-to-sit postural transition detection using pressure sensor arrays.

1.2 Motivation

There are good engineering and good clinical reasons to study recognition of sitting and lying postures and the transitions between them using an unobtrusive bed-based pressure sensor array. From the engineering perspective, recognizing postures and postural transitions present within a continuous stream of pressure data is useful because it enables meaningful segmentation of the pressure signal. Once the pressure signal is divided into segments representing both static postures and dynamic postural transitions, the segmented pressure information may be used as input to other signal processing systems which rely on previously segmented pressure data. For example, a system which estimates breathing rate from a person lying on the bed using pressure sensor arrays (such as [4]) might require the person to be lying on their back (supine) or stomach (prone) for the system to function correctly. As such, in a continuous patient monitoring scenario, a posture recognition system could be used to activate a breathing rate detector.

The majority of research into automated posture and postural transition detection is focused on the use of wearable sensors. While wearable sensors have the advantages of mobility and continuous daily monitoring, they also present a few practical drawbacks which are inherently avoided by using a bed-based pressure sensing system. The efficacy of a wearable sensor, such as an accelerometer, may be dependent on both its position and orientation when attached to the body [5]. In smart home environments where users may don wearable sensors with little or no help, the users may either forget to wear or incorrectly attach and orient wearable sensors. In contrast, a bed-based pressure sensor array could remain active, consistently in place, and properly oriented once installed. Portable devices, especially the smaller wearable ones, may also be misplaced. This situation is of no concern with a stationary bed-based pressure sensor array. Finally, a bed-based sensing system would likely be powered from a nearby outlet, whereas users of wireless wearable devices may have

to remember to charge their device batteries.

From the clinical perspective, the ability to provide information about a patients physical context is an important factor in comprehensive health monitoring [6]. The clinical usefulness of systems to unobtrusively detect posture and postural transitions also stems from the relevance of lying postures to a number of pathologies and from the use of postural transitions as indicators of mobility [7]. Sleep posture is a well known contributing factor in obstructive sleep apnea [8, 9, 10, 11, 12, 13, 14] and it has been shown that sleeping in a lateral posture (on one's side) reduces the severity of central sleep apnea [15]. The prone sleeping position has been associated with higher incidences of sudden infant death syndrome (SIDS) and is now considered an extrinsic risk factor for SIDS [16, 17, 18]. It has also been shown that symptomatic gastro-esophageal reflux can be effectively managed by sleeping in a left-lateral decubitus position [19, 20]. Combinations of posture recognition and biofeedback systems (e.g. [21]) can be used to detect undesirable lying positions in those suffering from these ailments and signal them to assume lying positions that promote their health.

The duration and quality of postural transitions, particularly sit-to-stand (SiSt) and lie-to-sit (LiSi), are useful indications of a person's mobility status. For example, these postural transitions are used in mobility assessment scales such as the Motor Assessment Scale (MAS) [22] and Mobility Scale for Acute Stroke Patients (MSAS) [23, 24] to assess the mobility of stroke patients in an acute setting. Automated and unobtrusive postural transition detectors could be used to continuously monitor changes in mobility to track rehabilitation progress and provide guidance in rehabilitation planning.

Overall, the ability to continuously monitor activities of daily living, including sleep and bed entry and exit, automatically, remotely, and unobtrusively would help alleviate the pressure placed on health care professionals. Encouragingly, these types of systems have in general been regarded positively by patients [25] and by seniors in

particular [26, 27].

1.3 Methodological objectives

To answer the first two parts of the research question, the feasibility of recognizing postures using bed-based optical pressure sensor arrays, the following methodological objectives were identified:

1. choose a set of relevant candidate pressure signal features;
2. choose a set of candidate pattern classification techniques;
3. determine the best set of candidate features for recognizing sitting and lying static activities;
4. compare the system's sitting and lying static activity recognition performance to results reported in the related literature;
5. determine the best set of candidate features for recognizing lying postures; and
6. compare the system's lying posture recognition performance to results reported in the related literature.

To answer the third part of the research question, the feasibility of detecting and recording the duration of lie-to-sit postural transitions using bed-based optical pressure sensor arrays, the following objectives were identified:

7. implement a pressure signal based movement detection algorithm;
8. design and implement a rule-based postural transition classifier which integrates movement detection and static activity recognition to detect lie-to-sit postural transitions;

9. determine an effective set of system parameters to achieve reliable and accurate lie-to-sit postural transition detection; and

10. compare the postural transition detector performance to an external standard.

1.4 Summary of contributions

Contributions provided by the preliminary work from which this thesis was based include:

1. demonstrating that a bed-based pressure sensor array can detect differences in postural transition duration between young healthy, older healthy, and older stroke patients;

    - [28] P. Carlson, N. Foubert, F. Knoefel, R. Goubran, M. Bilodeau, and H. Sveistrup, "Smart mat technology: Can it differentiate bed transfers in young healthy, old healthy and old stroke populations?" in *Canadian Geriatrics Society Annual General Meeting*, Toronto, Canada, April 2009

2. demonstrating that a bed-based pressure sensor array can detect differences in postural transition duration between older healthy and older hip fracture patients.

    - [29] S. Mondoux, P. Carlson, N. Foubert, A. Arcelus, F. Knoefel, R. Goubran, M. Bilodeau, and H. Sveistrup, "Smart mat technology: Can it differentiate bed transfers in older healthy and older hip fracture patients?" in *Regional Geriatrics Program Annual General Meeting*, Ottawa, Canada, October 2009.

In addition, this thesis contributes new knowledge by:

3. demonstrating, using simple pressure signal features, that bed-based optical pressure sensor arrays can reliably recognize lateral decubitus, prone, supine, and fetal lying postures when placed beneath a mattress typical of Canadian hospitals;

4. providing the first analysis of lie-to-sit postural transition detection using bed-based pressure sensor arrays;

5. demonstrating that bed-based pressure sensor arrays can be used to unobtrusively detect and determine the durations of lie-to-sit postural transitions with an accuracy commensurate of a medical student analyzing video recordings of the same postural transitions;

6. determining a set of simple pressure signal features such that bed-based *optical* pressure sensor arrays may be used to reliably recognize sitting and lying static activities in a variety of locations and orientations on a bed;

7. demonstrating that a nearest neighbour classifier is a better choice than a support vector machine for recognizing lying postures when limited training data is available; and

8. discovering that an effective fetal lying posture recognizer could be implemented using a nearest neighbour classification of a person's center of pressure on the mattress.

1.5 Thesis organization

This thesis is organized into six chapters in total. Chapter 2 discusses the background information required to understand the methods used in subsequent chapters and reviews relevant previously published work in pressure sensitive environments and physical context awareness. Chapter 3 reviews the experimental setup and data

collection methods used to gather the data required to evaluate the performance of the posture recognizers and the postural transition detector. Chapter 4 presents the proposed static activity and lying posture recognition systems and reports results of performance evaluation for these systems. Chapter 5 presents the proposed postural transition detector and reports results of its performance evaluation. The final chapter presents the conclusions of the thesis and examines potential avenues of future research to extend the results of this work.

Chapter 2

Background and Related Work

The static activity and lying posture recognizers use common pattern recognition techniques. This chapter begins with a brief overview of pattern recognition, focusing on support vector machines (SVMs) and nearest neighbour (NN) classifiers as they are used to implement the pattern recognizers in this thesis. This is followed by a discussion of the metrics and techniques used to evaluate the performance of pattern recognizers. The last sections review related work in both pressure sensitive home environments and physical context awareness.

2.1 Pattern recognition

The goal of a pattern recognition system is to learn a target concept (i.e. pattern) by analyzing a set of features extracted from raw data, usually with the intention of taking an action or making a decision based on the type of pattern recognized [30]. There are three basic pattern learning methods: unsupervised, supervised, and reinforcement. Unsupervised learning occurs when a pattern recognition system must learn a concept from a data set without being given any prior information about the structure of the concept to be learned. In supervised learning, the system is provided with a number of examples of the concept from which it attempts to learn the correct

structure. With reinforcement learning, the learner is not given correct examples, but instead is rewarded for correctly recognizing patterns. The methods used in this thesis all fall under the supervised learning paradigm.

The basic problem of supervised learning in the simplest case of discriminating between two classes may be expressed in the following way. Given a set of m training data instances (the **training set**) $\{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_m, y_m)\}$, where the $\mathbf{x}_i \in \mathbb{R}^F$ are feature vectors (combining F features) which describe the data instances and the $y_i \in \{-1, 1\}$ are the instance classes, the supervised learning problem is to infer a mapping

$$f : \mathbb{R}^F \rightarrow \{-1, 1\}$$

from the training set which correctly maps new feature vectors to classes. The learner is then tasked with classifying new data instances (e.g. from a **test set**) using this mapping. Numerous supervised learning methods have been invented including perceptrons, neural networks, decision trees, Bayes nets, naive Bayes, support vector machines, nearest neighbour, and many more [30].

2.1.1 Support vector machine classification

SVMs belong to the class of discriminant methods called maximum-margin classifiers, in which the key process is finding a separating hyperplane which maximizes the margin on either side of the hyperplane separating two target classes in the feature space [31]. A discriminant method does not attempt to estimate the parameters of the underlying probability densities of the data being modeled; rather, it directly determines the parameters of a function which discriminates (i.e. separates) the classes in feature space [30].

The following brief mathematical description of SVMs is derived from [31]. SVMs are capable of finding non-linear discriminant surfaces by employing the so-called

kernel trick whereby a (kernel) function is used to map the feature vectors to a higher dimensional space. In the higher dimensional space, the SVM uses an optimization procedure to determine the hyperplane which best separates the classes using the maximum margin criterion. The linear discriminant surface is then transformed back into the original feature space, often resulting in a non-linear discriminant surface in the original feature space.

To find the maximum-margin linear discriminant surface the following optimization problem is solved:

$$\min_{\mathbf{w},b,\xi} \quad \frac{1}{2}\mathbf{w}^T\mathbf{w} + C\sum_{i=1}^{l}\xi_i \tag{2.1}$$

$$\text{subject to} \quad y_i(\mathbf{w}^T\phi(\mathbf{x}_i) + b) \geq 1 - \xi_i, \tag{2.2}$$

$$\xi_i \geq 0 \tag{2.3}$$

where $\mathbf{w}$ is the normal vector of the separating hyperplane, ξ_i is an error tolerance, C is a penalty parameter, b is the bias, and ϕ is a function which maps the feature vectors $\mathbf{x}_i$ into a higher dimensional space (recall: ϕ may be a non-linear map). The optimization problem is typically solved using specialized algorithms such as sequential minimal optimization [32].

A kernel function, or kernel, is defined by

$$K(\mathbf{x}_i, \mathbf{x}_j) \equiv \phi(\mathbf{x}_i)^T\phi(\mathbf{x}_j) \tag{2.4}$$

There are several kernels commonly used in SVM applications, such as linear, polynomial, sigmoid, and radial basis function (RBF) [33]. A Gaussian function is often used as the RBF, and in such cases the RBF kernel takes the form

$$K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma||\mathbf{x}_i - \mathbf{x}_j||^2). \tag{2.5}$$

Using an RBF kernel, the SVM has two hyperparameters which must be specified before training. The C hyperparameter is a misclassification penalty which appears in the specification of the SVM optimization problem (equation 2.1). The γ hyperparameter controls the extent of the RBF and therefore controls the smoothness of the discriminant surface.

For multi-class classification, the SVM classifier can use a set of $k(k-1)/2$ binary classifiers for a k class problem, and the class is chosen by majority vote among the binary classifiers.

SVMs are popular and have been used for many other classification problems in biomedical engineering including automated gait classification [34], classification of limb motion for myoelectric control [35], automatic detection of voice pathologies [36], identification of retinal arteries and veins [37], heartbeat recognition [38], and physical activity recognition using accelerometer data [39], to name a few.

2.1.2 Nearest neighbour classification

The nearest neighbour (NN) algorithm is one of the simplest classification techniques. Given a feature vector $\mathbf{x}$ to classify, it simply chooses the class y_i of the nearest neighbour $\mathbf{x}_i$ from a training set $\{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_m, y_m)\}$, where the nearest neighbour is determined using a particular distance metric in the feature space (e.g. Euclidean distance) [40]. Similarly, the k-nearest neighbour (kNN) classifier chooses a class based on the k nearest neighbours in the feature space.

In this thesis, the k nearest neighbours are determined using the standard Euclidean distance between feature vectors $\mathbf{a} = [a_1, a_2, \ldots, a_m]^T$ and $\mathbf{b} = [b_1, b_2, \ldots, b_m]^T$

$$d(\mathbf{a}, \mathbf{b}) = \sqrt{\sum_{i=1}^{m} (a_i - b_i)^2} \tag{2.6}$$

The NN and kNN classifiers also use a weighted class voting scheme, with the weight

of a neighbour computed as $\exp(-t^2/s^2)$ where t is the rank of the neighbour in the list of k neighbours ordered by Euclidean distance from the feature vector in question and s is a constant scale factor. This weighting scheme ensures that the nearer neighbours have more influence on the classification decision; the constant s is chosen such that the weight of the furthest neighbour is 0.001 [41].

NN and kNN classifiers have several advantages including the ability to learn from small training sets, simple on-line learning, and they often perform competitively with more sophisticated classification methods [42]. Furthermore, they have been applied in various biomedical engineering studies including classification of knee gait patterns [43], epileptic seizure detection [44], and postural transition detection [45].

2.1.3 Performance evaluation

Evaluating the performance of pattern recognition systems typically involves training the system with one set of example data instances in which the true classes are provided and then testing the system on a separate set of instances in which the classes are hidden or unknown. There are two common modes of performance evaluation used in the activity and posture recognition literature. In subject-dependent analysis both the training set and test set contain data from the same individuals. In subject-independent analysis the test set contains data from subjects for which no training data is made available.

Since it is generally impossible to train a pattern recognizer on all the patterns it will have to recognize in the future, performance evaluation often consists of estimating the recognizers ability to generalize its classification capability by using mutually exclusive training and test sets. Often, however, only a relatively little amount of data is available for training and testing. In such cases, data resampling techniques are used to reduce the impact of this limitation. Two common resampling techniques described in the relevant literature are k-fold cross-validation and leave-one-subject-out

validation.

In k-fold cross-validation, the set of all available data is split into k mutually exclusive subsets. The recognition system is then trained using the data belonging to $k-1$ of the subsets and is tested on the remaining subset. This process is repeated for each of the k possible test sets, resulting in k performance metrics which are then averaged to produce a final estimated performance metric.

Leave-one-subject-out validation is a specific type of cross-validation in which the data for all but one subject is used for training and data from the left-out subject is used for testing. Again, this process is repeated for all subjects to produce a single averaged performance metric.

These resampling techniques are general in that they do not require using any particular performance metric. There are a number of common performance metrics used to characterize pattern recognizers. The classes in a two-class (i.e. binary) classification task are commonly referred to as the positive class and the negative class, and there are four possible outcomes when a binary classifier is applied to a particular data instance:

- **True positive (TP),** e.g. sitting correctly recognized as sitting;
- **False positive (FP),** e.g. lying incorrectly recognized as sitting;
- **True negative (TN),** e.g. lying correctly recognized as lying; and
- **False negative (FN)**, e.g. sitting incorrectly recognized as lying.

Using these definitions, the following list defines and describes some of the most commonly used classification performance metrics:

- **Classification accuracy** is the percentage of classes which are correctly rec-

ognized by the classifier, defined as

$$\text{accuracy} = \frac{\text{number of TPs} + \text{number of TNs}}{\text{number of TPs} + \text{number of TNs} + \text{number of FPs} + \text{number of FNs}} \tag{2.7}$$

In the case of multi-class recognition, this definition is generalized as

$$\text{accuracy} = \frac{\#\text{ correctly classified}}{\#} \tag{2.8}$$

where # means 'number of instances'.

- **Sensitivity** is a measure of how well the classifier avoids false negatives, defined as

$$\text{sensitivity} = \frac{\text{number of TPs}}{\text{number of TPs} + \text{number of FNs}} \tag{2.9}$$

In the multi-class case sensitivity is defined with respect to a particular class τ and is generalized as

$$\text{sensitivity}(\tau) = \frac{\#\text{ correctly classified as } \tau}{\#\text{ correctly classified as } \tau + \#\text{ misclassified as not } \tau} \tag{2.10}$$

where τ is the class of interest.

- **Specificity** is a measure of how well the classifier avoids false positives, defined as

$$\text{specificity} = \frac{\text{number of TNs}}{\text{number of TNs} + \text{number of FPs}} \tag{2.11}$$

In the multi-class case specificity is defined with respect to a particular class τ and is generalized as

$$\text{speficity}(\tau) = \frac{\#\text{ correctly classified as not } \tau}{\#\text{ correctly classified as not } \tau + \#\text{ instances misclassified as } \tau} \tag{2.12}$$

- **Precision** is defined as

$$\text{precision} = \frac{\text{number of TPs}}{\text{number of TPs} + \text{number of FPs}} \tag{2.13}$$

In the multi-class case precision is defined with respect to a particular class τ and is generalized as

$$\text{precision}(\tau) = \frac{\#\text{ correctly classified as } \tau}{\#\text{ correctly classified as } \tau + \#\text{ misclassified as } \tau} \tag{2.14}$$

- **Recall** is defined as

$$\text{recall} = \frac{\text{number of TPs}}{\text{number of TPs} + \text{number of FNs}} \tag{2.15}$$

In the multi-class case recall is defined with respect to a particular class τ and is generalized as

$$\text{recall}(\tau) = \frac{\#\text{ correctly classified as } \tau}{\#\text{ correctly classified as } \tau + \#\text{ misclassified as not } \tau} \tag{2.16}$$

2.2 Pressure Sensitive Environments

A pressure sensitive environment uses pressure sensors to collect data representing the tactile interactions between occupants and their environment [3]. The research presented herein belongs to a family of related studies in pressure sensitive environments being conducted at Carleton University as part of the Technology Assisted Friendly Environments for the Third Age (TAFETA): Smart Systems for Health program of research [46]. Within this program of research, pressure sensitive mats have been used to detect respiratory rate [47, 4], movement onset [48], and rollover [49]. Pressure sensitive mats have also been used to detect simulated central sleep apnea using both

single-sensor and multi-sensor voting methods [50]. Arcelus et al. [51] outfitted commode grab bars with pressure sensors at specified locations to extract various clinical parameters including maximum force and contact location to characterize commode sit-to-stand and stand-to-sit body weight transfers.

Other research groups have investigated similar applications of pressure sensitive technology. Tan et al. [52], Slivovsky and Tan [53] implemented a sitting posture classification system using a pressure sensing chair. Pressure sensitive floors have been used for a variety of reasons including determining the locations of people or objects [54, 55], identifying individuals [56, 57, 58, 59], and performing gait analysis [60].

2.3 Physical Context Awareness

Two important functions of a smart home environment for health care applications are to assist the occupants in their activities of daily living and to monitor their health status. In order to accomplish these tasks most effectively, a smart environment should be capable of recognizing the occupants physical context in order to apply the appropriate assistance and monitoring procedures [61]. A smart environment with context awareness can detect changes in the state of its occupants and adapt to such changes [62].

Using pressure sensitive mats, this thesis focuses on the study of three inter-related context recognition scenarios: 1) recognizing when a person is sitting at the edge of the bed or lying in the bed (i.e. static activities), 2) when the person is lying in bed, recognizing the person's lying posture, and 3) detecting when a person transitions from lying in bed to sitting in bed. The following sections highlight related work in these three areas.

2.3.1 Recognizing Static Activities

One aspect of physical context detectable using sensors is physical activity. Detecting physical activity can provide essential context information for health care monitoring applications. For example, a heart rate monitor need not raise the alarm for elevated heart rate if the system detects that its user is currently exercising. In the research on physical activity recognition, activities are typically categorized as being either static (e.g. sitting or lying) or dynamic (e.g. walking). This thesis focuses exclusively on static activity recognition.

An active area of research in physical activity recognition is centered on the use of simple wearable sensors. Numerous sensor configurations have been studied with varying success including a single triaxial accelerometer [63, 64], multiple triaxial accelerometers [65], multiple biaxial accelerometers [66, 67], a combination of five wireless triaxial accelerometers and a heart-rate monitor [68], and a combination of one biaxial accelerometer and a gyroscope [69]. These studies typically investigate both static and dynamic activities, but they report static activity recognition accuracies in the range of 88-100% for subject-dependent tests and 84-99% for subject independent tests.

Pressure sensitive surfaces have previously been used to detect physical activities. Pressure sensitive floors have been used to analyze gait [70] and dance movements [71]. Tan et al. [52] used a pressure sensitive chair to classify a persons sitting posture. Pressure sensitive mats have also been used to classify the lying postures of users in bed, as discussed further in the next section.

2.3.2 Lying Posture Recognition

In much of the physical activity recognition literature the classification of static activities often goes no deeper than basic categories such as 'sitting' or 'lying'. However, there has been some research in the area of using wearable sensors to recognize sleep

(i.e. lying) postures. Baker et al. [72] developed a number of in-home health care sensors, including an accelerometer-based wearable sensor for detecting sleep position in infants. Van Laerhoven et al. [73] used a wrist-worn sensor including a 3-axis accelerometer and nine tilt switches to determine adult sleep position and claimed an average precision of 80% for discriminating between left/right lateral, prone, and supine lying postures. A head mounted accelerometer was used in [21] to detect head position during sleep for providing biofeedback when patients with obstructive sleep apnea assume the supine position.

Some previous research has investigated lying posture classification using pressure sensor arrays. Hsia et al. [74] built a custom sensor array using 16 long and narrow force sensing resistors (FSRs), computing skewness and kurtosis in the pressure distribution to detect supine and left/right lateral postures. They claim to have achieved 100% accuracy when the subjects' bodies were laid parallel to the mat center line and 78.8% average accuracy. Seo et al. [75] claimed a 93.6% successful recognition rate using 336 FSRs embedded in the top surface of a mattress with principal component analysis (PCA) and a radial-basis function (RBF) neural network to detect supine, left/right lateral, and sitting postures. Harada et al. [76] placed a pressure sensor array consisting of 210 FSRs underneath a 5 cm thick futon-mat to detect supine and left/right lateral postures, and figures in their paper suggest that they achieved 100% accuracy. Finally, Nishida et al. [77] placed an array of 221 FSRs beneath a 10 cm thick futon-mat to detect supine, prone, and left/right lateral postures and evaluated their system by comparing the number of posture changes detected by the system with those detected by video analysis.

2.3.3 Postural Transition Detection

Sit-to-stand (SiSt) and stand-to-sit (StSi) postural transition detection has received much attention in the literature. The detection of these postural transitions has

been investigated using a number of sensor modalities and methods. Arcelus et al. [78] used pressure sensitive mats placed underneath a bed mattress and on the floor adjacent to the bed to record the duration of manually segmented SiSt transitions using center of pressure and the discrete wavelet transform (DWT). Najafi et al. [79] used a single kinematic sensor, comprised of one gyroscope and two uniaxial accelerometers, attached to the chest to detect SiSt and StSi postural transitions during continuous monitoring of a persons physical activity. Jafari et al. [45] used a three-axis accelerometer for detecting SiSt, StSi, stand-to-lie (StLi), and lie-to-stand (LiSt) transitions using k-nearest neighbour and support vector machine classifiers. A single three-axis accelerometer, mounted at the waist, was used by Mathie et al. [80] to detect SiSt and StSi transitions by thresholding the acceleration signal magnitude area and by Allen et al. [81] to detect SiSt, StSi, LiSt, and StLi transitions using both rule-based and Gaussian mixture model classifiers. SiSt and StSi durations were measured via waist-worn triaxial accelerometer using a DWT on the acceleration signal vector magnitude in [82]. Costantini et al. [83] used a custom sensor composed of three accelerometers and three gyroscopes to measure StSi and SiSt durations via spectral analysis. In [69] the DWT was used with a chest-mounted gyroscope to determine SiSt and StSi durations. Salarian et al. [84] employed a fuzzy classifier to classify and measure the duration of SiSt and StSi transitions using one biaxial accelerometer and two gyroscopes. A markerless computer vision system was implemented in [85] to determine SiSt transition duration.

Other modalities have been used to characterize SiSt and StSi transitions mechanically. Piezo-electric force-plates, video, and electromyographic data were used in [86] to characterize SiSt transitions in a population of healthy elderly subjects. Force-plates were used in [87] to identify features of the SiSt and StSi transitions in stroke patients who are at risk of falling. Riley et al. [88] used video kinematic analysis to differentiate the mechanisms in failed SiSt transitions from successful SiSt

transitions. Video kinematic analysis was used in [89] to analyze whole-body motion during the SiSt postural transition in young healthy subjects. A markerless video system was investigated in [85] to detect SiSt in young and elderly people.

Compared with SiSt and StSi postural transitions, there is a relative dearth of published research investigating the detection of LiSi and SiLi transitions. Yang et al. [90] describe a wearable motion sensor network consisting of eight sensors, each integrating a triaxial accelerometer and a biaxial gyroscope, for recognition of human actions including LiSi, SiLi, StSi, and SiSt postural transitions. Their system uses a distributed action recognition algorithm for the simultaneous segmentation and classification of actions from the sensor signals. In [91] the authors employed the same sensor configuration but used a standard-deviation based signal segmentation routine to recognize human activity including SiLi and LiSi transitions. Finally, Najafi et al. [79] detected postural transitions including SiLi and LiSi using kinematic sensors integrating two uniaxial accelerometers and a gyroscope and the DWT.

Chapter 3

Experiment Setup and Data Collection

The methodology used to evaluate the posture recognition and postural transition detection systems described in the forthcoming chapters used four different data sets. To clarify which data set is being discussed, each set was given a particular name; the four data sets are known as the **subject-dependent static activity data**, **subject-independent static activity data**, **lying posture data**, and **postural transition data**.

Furthermore, these data sets were collected in two different environments, which are referred to as the **clinical environment** and the **laboratory environment**. The subject-independent and postural transition data were collected in the clinical

Data set	Collection environment	Used in chapter
Subject-dependent static activity	Laboratory	4
Subject-independent static activity	Clinical	4
Lying posture	Laboratory	4
Postural transition	Clinical	5

Table 3.1: Summary of data sets used for performance evaluation.

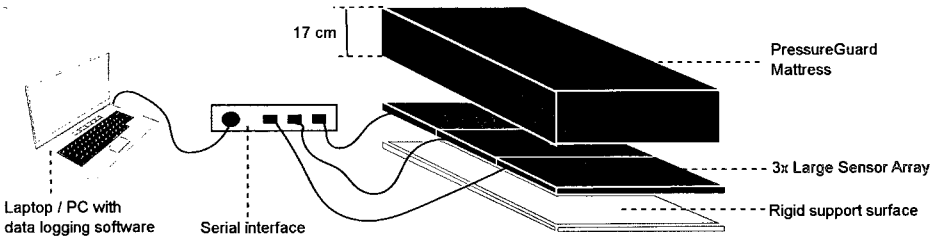

Figure 3.1: Pressure data collection equipment setup.

environment, while the subject-dependent data and lying posture data were collected in the laboratory environment. The clinical environment was located at the Elisabeth Bruyere Hospital in Ottawa, Ontario and the laboratory environment was located at the Digital Signal Processing lab at Carleton University in Ottawa, Ontario. The data sets, environments, and chapters in which the data sets were used are summarized in table 3.1.

3.1 Equipment setup

Figure 3.1 depicts the basic data collection equipment setup. The equipment setup consisted of six pressure sensitive mats, placed side-by-side, on top of a rigid support surface and underneath a mattress. In the clinical environment, the support surface was a typical hospital bed frame. In the laboratory environment, the surface was a long table. The six pressure sensitive mats are integrated into three physical mats, as each physical mat is composed of two pressure sensitive component mats.

The pressure sensitive mats were fibre-optic based Large Sensor Arrays provided by Tactex Controls Inc [92]. Each mat is approximately 66 cm in width and 71 cm in height and consists of 44 Kinotex® sensors positioned between layers of foam. Kinotex sensors leverage the pressure dependent light scattering properties of the foam material to detect forces applied to the foam [93]. All sensors are sampled at a frequency of 10 Hz and are spaced in a honeycomb tessellation with an inter-sensor distance of approximately 13 cm. Figure 3.2 depicts the sensor locations when the

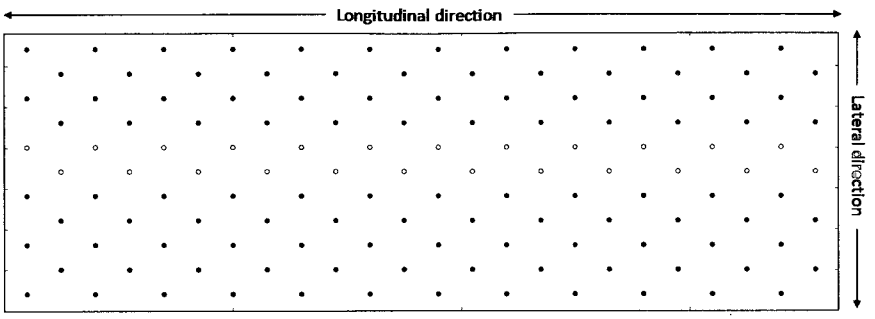

Figure 3.2: Sensor locations when all six sensor groups (three physical mats) are placed side-by-side and considered a single sensor array

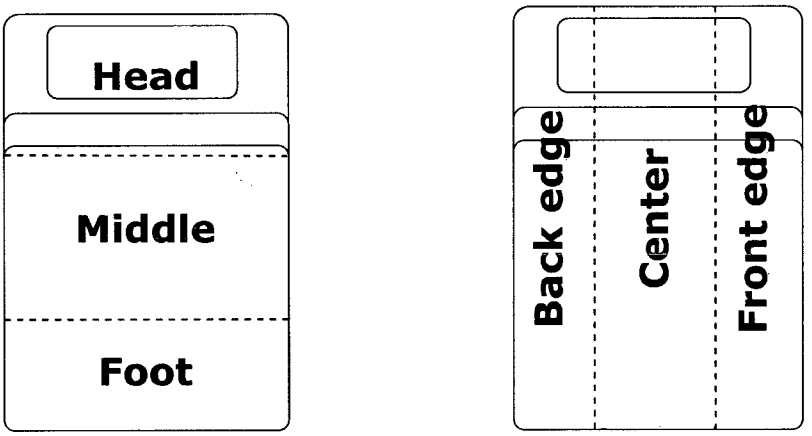

(a) Longitudinal divisions of the mattress.

(b) Lateral divisions of the mattress.

Figure 3.3: Longitudinal and lateral divisions of mattress surface area used to describe location of sitting and lying activities.

three physical mats are placed side-by-side.

Each physical mat is composed of two sensor groups with their own control circuits. Each sensor group was connected to a RS-232 interface using a DB9-RS232 cable and RJ-45 connector. All three physical mats (6 connector cables in total) were connected to a serial adapter which in turn is connected to a serial port (RS-232) on the data logging computer. The computer was running data logging software provided by Tactex. Note that figure 3.1 shows all RS-232 interfaces and the adapter as a single unit for the sake of simplicity.

The mattress was a PressureGuard Renew™, manufactured by Span-America Medical Systems Inc. [94], measuring 17 cm in thickness. The PressureGuard Renew is

specifically designed to evenly distribute pressure to promote patient comfort, which in effect reduces the capability for the underlying pressure sensor array to resolve details such as limb placement on the mattress.

Figure 3.3 depicts the spatial divisions of the mattress surface used in the forthcoming discussion. The mattress is divided into three regions in both the longitudinal and lateral directions. The longitudinal regions (figure 3.3a) include the head, middle, and foot regions. The lateral regions (figure 3.3b) include the back edge, center, and front edge. In both the clinical and laboratory environments, the back edge was against a wall, so study participants always exited the bed from the front edge.

In addition to the pressure sensitive mats, two Sony high-definition video cameras were used to record participant activities in the clinical environment.

The laboratory environment equipment setup was designed to mirror the clinical environment setup as closely as possible. The pressure sensitive mats, serial interfaces, cables, data logging software, and mattress were identical. As there was no hospital bed frame available in the laboratory, the mattress and pressure sensitive mats were instead placed on a long table. Other differences were the absence of video cameras (video data was not required) and the pressure sensitive mats were connected to a desktop PC instead of a laptop.

3.2 Data format

Time stamp	Sample number	Event number	Pressure	Sensor 1 Value	Sensor 2 Value	...	Sensor 132 Value

Table 3.2: Data format output by Tactex data logging software

The data logging software recorded data in the format displayed in table 3.2. Each row of the data file contained a time stamp, incremental sample number, event number, pressure value, and the value of all 132 sensors. A change in event number indicated a manual mark point created by the user of the data logging software (i.e. an 'event'

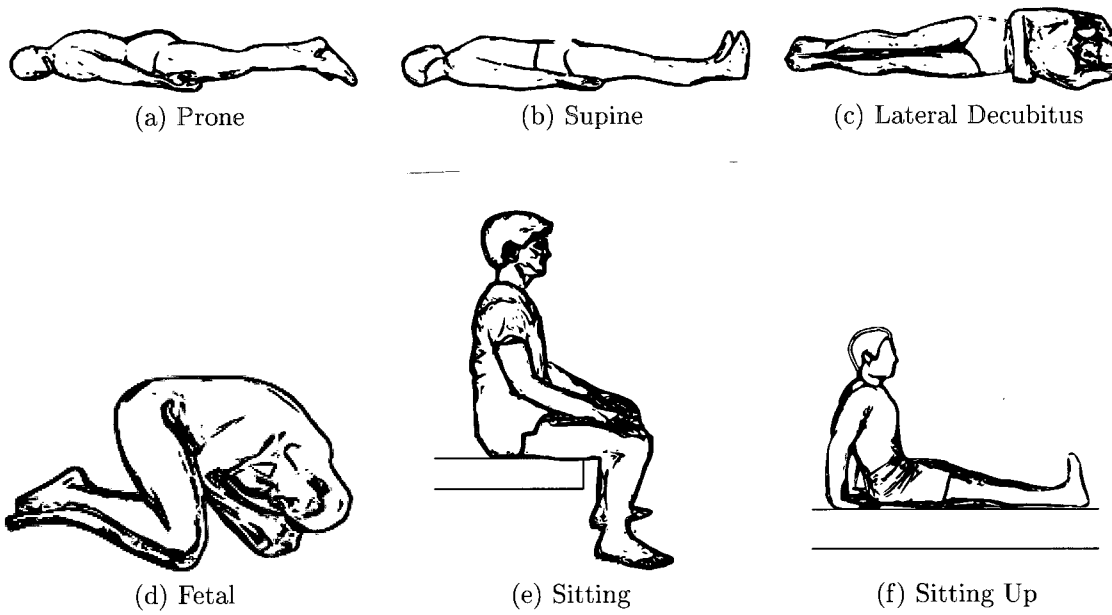

(a) Prone (b) Supine (c) Lateral Decubitus

(d) Fetal (e) Sitting (f) Sitting Up

Figure 3.4: Postures considered in this thesis.

was logged). Of all the data fields, only the event number and the 132 sensor value columns were used.

3.3 Postures

Figure 3.4 depicts the set of postures considered in this thesis. The person is in a **prone** lying posture (figure 3.4a) when lying on his stomach and in a **supine** lying posture (figure 3.4b) when lying on his back. Lying on the side of the body is called a **lateral decubitus** lying posture (figure 3.4c). A **fetal** lying posture (figure 3.4d) is a variation on the lateral decubitus posture in which the knees are brought up towards the chest. Both lateral decubitus (hereafter simply referred to as lateral) and fetal postures are further specified by the side of the body which touches the mattress (e.g. left lateral). This thesis discriminates between two different types of sitting postures. The **sitting** posture (figure 3.4e) refers to when the person is sitting at the edge of the mattress with his feet either dangling over the side of the bed or placed on the floor adjacent to the bed. Finally, when the person is sitting up in bed with legs

extended in the longitudinal direction, he is in a **sitting up** posture (figure 3.4f) .

3.4 Collection protocol

As the four data sets were used to evaluate different parts of the posture recognizer and postural transition detector, the data sets were collected using different protocols. This section reviews the collection protocols used to collect the subject-dependent and subject-independent static activity data, lying posture data, and postural transition data.

3.4.1 Static activity data

The static activity data consists of two separate data sets: the subject-dependent data set and the subject-independent data set. These data sets are so named to reflect the type of evaluation method for which they are used (i.e. subject-dependent and subject-independent).

Subject-dependent data set

The subject-dependent data was captured in the laboratory environment, and consisted of a series of performances of sitting and lying activities by one 27 year old healthy male. This data was used to evaluate the ability of the static activity recognizer to train on a person and then recognize sitting and lying activities of the same person. The location and orientation of each activity was varied with respect to the mattress, anticipating realistic real-world scenarios.

Each sitting activity was specified by three qualifiers: the longitudinal position, lateral position, and hand position. The three possible longitudinal positions and three possible lateral positions corresponded to the longitudinal and lateral regions as described in section 3.1 and depicted in figure 3.3. The three possible hand positions

included two hands at side, one hand at side, and no hands at side (e.g. placed on the knees). Note that in this data set, the sitting posture was always 'sitting' as depicted in figure 3.4e and *not* 'sitting up' as depicted in figure 3.4f.

Each lying activity was also specified by three qualifiers: the lying posture, orientation, and lateral position. Six lying postures were considered (see figure 3.4): prone, supine, left lateral, right lateral, left fetal, and right fetal. The orientation was either head-to-head or head-to-foot; head-to-head means that the persons head was located at the head of the bed (i.e. the standard location of a pillow) and their feet were located towards the foot of the bed, while head-to-foot describes the opposite (i.e. the persons head was located at the foot of the bed). The lateral position referred to the lateral location of the person relative to the mattress. The lateral positions considered included back edge, middle, and forward edge (see figure 3.3b). Note that the lateral position was only specified for the lateral decubitus postures because the other postures typically cover more than one lateral position.

The full range of sitting and lying activities (i.e. comprising all possible combinations of posture, orientation, position, etc.) were performed by the person (47 in total) and each activity performance was marked using the data logging software to enable manual segmentation and classification.

Subject-independent data set

The subject-independent data was captured in the clinical environment. This data was collected from 25 volunteer participants representing four populations: young healthy, older healthy, older post-stroke, and older post-hip-fracture. A summary of participant statistics is reported in table 3.3. The older post-stroke and older post-hip-fracture groups are hereafter referred to simply as the stroke and hip fracture groups, respectively. Each participant was asked to perform a series of bed entry and exit routines. However, prior to performance of bed entry and exit routines, the

Participant Group	Number of Participants	Mean Height (cm)	Mean Weight (kg)
Young Healthy	10	162.3 ± 5.1	52.3 ± 13.9
Older Healthy	5	164.6 ± 9.5	67.2 ± 12.9
Hip Fracture	5	160.2 ± 17.1	92.5 ± 21.1
Stroke	5	172.7 ± 9.2	58.9 ± 14.8

Table 3.3: Statistics of participants in the clinical environment.

participants were asked to both sit and lie (supine), relaxed, on the bed to collect baseline data for a short period of time. This baseline relaxed sitting and lying data was used for subject-independent analysis. The participant data was manually segmented using the data logging software by a research assistant during collection, enabling later manual posture classification. Unlike with the subject-dependent data, the sitting and lying positions were restricted to head-to-head supine lying positions and center front edge sitting positions, as these were the positions used for bed entry and exit routines.

3.4.2 Lying posture data

The lying posture data was collected in the laboratory environment and consists of a series of different lying postures. As with the subject-dependent static activity data, all the lying posture data consists of performances by a one 27 year old healthy male. In contrast to the subject-dependent static activity data, all the lying postures were performed in a head-to-head orientation with the body positioned over the center line of the mattress. The rationale for restricting orientation in this manner is that head-to-head, body over center line positioning corresponds to the most common orientation of people in bed. The lying postures included: supine, prone, left lateral, right lateral, left fetal, right fetal, and sitting up (see figure 3.4). The initiation time for each posture performance was marked using the data logging software to enable later manual segmentation and classification of the lying postures.

3.4.3 Postural transition data

The postural transition data was collected in the clinical environment from the same set of participants as in the subject-independent static activity data, summarized in table 3.3. Each participant performed ten bed entry and exit routines, with each routine composed of the following sequence of activities: sitting, standing, sitting, lying (supine), sitting. Participants were video recorded by two video cameras during the routine performances. One video camera was placed at the foot of the bed with the camera lens principal axis approximately co-linear to the mattress center line, the other camera was placed with its principal axis orthogonal to the principal axis of the first camera (i.e. facing the side of the bed).

The video data was later analyzed by two medical students to visually determine lie-to-sit postural transition durations and the lie-to-sit duration for each bed entry and exit performance for each participant was recorded in a spreadsheet. The pressure data was saved for later off-line analysis.

Chapter 4

Posture Recognition

The goal of the posture recognition system presented in this thesis is to unobtrusively and accurately detect the posture of a subject resting on a bed, using an optical pressure sensor array placed beneath the mattress. In the context of this thesis, posture recognition comprises two distinct processes: static activity (i.e. sitting or lying) recognition and lying posture recognition.

Figure 4.1 depicts a diagram of the the information flow from pressure sensor array through to posture recognition. The input to either the static activity recognizer or lying posture recognizer is unprocessed pressure data output from the pressure sensor array. The pressure data is passed into the static activity recognizer which determines if the person on the bed is sitting or lying. If the person is lying, the lying posture recognizer then determines the lying posture.

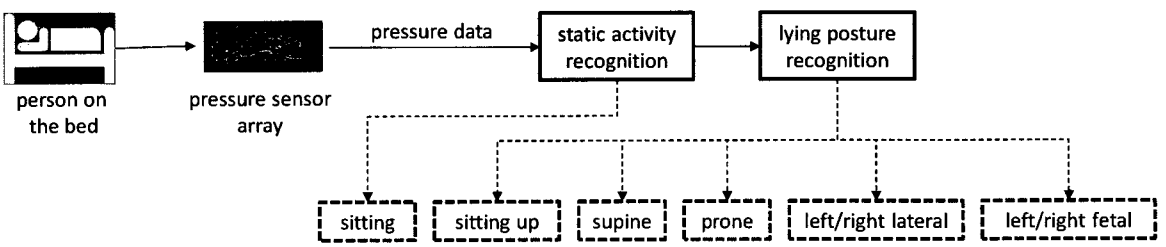

Figure 4.1: Information flow from the person on the bed to static activity and lying posture recognition.

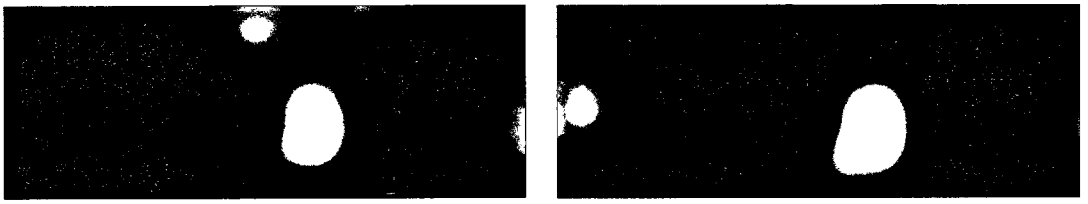

(a) Sitting with both hands down (b) Lying in left lateral posture

Figure 4.2: Comparing pressure distributions of sitting and lying postures. These images represent interpolated versions of the true sensor data (generated by convolution with a Gaussian).

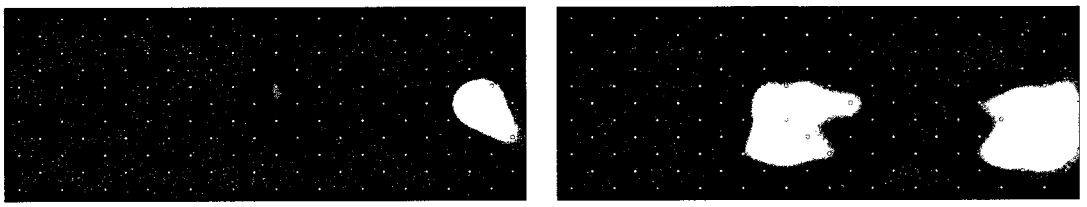

(a) No mattress (b) With mattress

Figure 4.3: Comparing pressure distributions with and without the mattress. These images represent interpolated versions of the true sensor data (generated by convolution with a Gaussian). The white dots show the sensor locations.

It might seem straightforward to discriminate between sitting and lying; after all, we expect that when a person is lying they span a much larger surface area of the mattress compared with sitting. Figure 4.2 shows why this is not always the case. Figure 4.2a shows an interpolated view of the pressure distribution with a person sitting in the middle of the bed with both hands down on the mattress and figure 4.2b shows the same person lying in the center of the bed in the left lateral posture. The hospital mattress has a smearing effect on the pressure distribution as it is recorded by the pressure sensor array. Because of this smearing effect–and the fact that the pressure sensors are already 13 cm apart–particular sitting postures may not be easily distinguished from particular lying postures. This smearing effect is demonstrated in figure X, which shows a left-lateral lying posture when the person was lying directly on the pressure sensor array and the same posture when the person was lying on the mattress placed on top of the pressure sensor array.

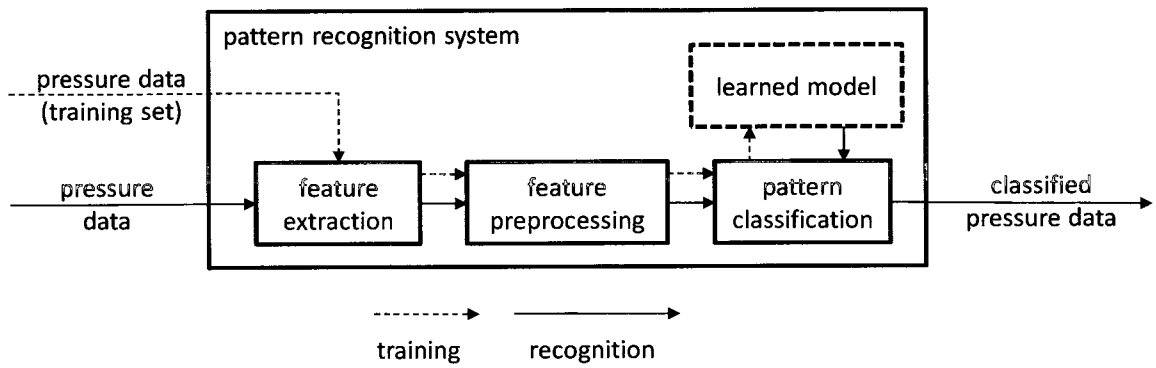

Figure 4.4: A pattern recognition framework for classifying pressure data.

4.1 Pattern recognition methodology

Both the static activity recognizer and lying posture recognizer follow the canonical design of many pattern recognition systems, shown in figure 4.4. From the unprocessed pressure data, relevant pressure signal features are extracted, preprocessed, and passed to a pattern classifier which has been previously trained with correctly classified patterns.

4.1.1 Feature extraction

Due to the low spatial resolution provided by the pressure sensor array in combination with the hospital mattress, more sophisticated geometry based feature extraction techniques (such as edge/shape extraction) were ruled out. Figure 4.5 shows an interpolated pressure distribution image of a person assuming a supine, prone, and left lateral posture. These figures qualitatively demonstrate the lack of informative shape information in the pressure distributions. Morphological techniques such as blob counting were also ruled out because there did not appear to be a consistent difference in the number of active regions when comparing sitting versus lying postures. Intuitively we might expect fewer active regions when a person is sitting compared to when they are lying, but figure 4.2 shows this is not always the case. Since geometrical and morphological features were unlikely to produce desirable results, simple

(a) Supine (b) Prone (c) Left lateral

Figure 4.5: Pressure distribution generated by three different lying postures. These images represent interpolated versions of the true sensor data (generated by convolution with a Gaussian). The white dots show the sensor locations.

Feature	Shorthand	Equation
(1) Sum of sensor values	SSV	4.1
(2) Number of active sensors	NAS	4.2
(3) Weighted sum of sensor values	WSSV	4.4
(4) Weighted number of active sensors	WNAS	4.6
(5) Longitudinal center of pressure	LonCP	4.7
(6) Lateral center of pressure	LatCP	4.8
(7) Longitudinal variance	LonV	4.9
(8) Lateral variance	LatV	4.10

Table 4.1: Summary of candidate pressure signal features.

descriptive statistics were chosen as candidate features.

In the following discussion, the pressure sensor array is represented as a set of indexed sensors $\{S_1[t], S_2[t], \ldots, S_N[t]\}$ where $N = 132$ is the total number of sensors in the array (see section 3.1). Each sensor is represented as a triple $S_i[t] = (x_i, y_i, s_i[t])$, where x_i, y_i are the x (longitudinal) and y (lateral) coordinates of the ith sensor on the sensor array surface, and $s_i[t]$ is the sensor value at time t. Using this notation, eight candidate sensor signal features are proposed and summarized in table 4.1.

(1) Sum of sensor values (SSV)

The sum of sensor values provides a measure of the total pressure exerted on the surface of the pressure sensor array. There is an intuitive argument for considering this a relevant feature for discriminating between sitting and lying. When someone sits on a bed with their legs over the side, a fraction of their body weight may be

borne by the floor and not by the bed. It follows that the pressure exerted on the sensor array by a person sitting at the bed's edge may be less than the pressure exerted when the same person is lying, assuming that while lying all limbs are resting on the bed.

The sum of sensor values at time t is a summation of the pressure value at t for all N sensors. This feature is defined mathematically as

$$\text{SSV}[t] = \sum_{i=1}^{N} s_i[t]. \tag{4.1}$$

(2) Number of active sensors (NAS)

The number of active sensors is another simple feature to compute. It provides a measure of the spatial distribution of a person's body on the pressure sensor array. Intuitively, the surface area of the mattress covered by a person while lying is expected to be typically greater than the area covered by the same person while sitting, so this feature may be useful for the discriminating between sitting and lying postures.

The number of active sensors at time t is the number of all sensors with pressure value greater than zero. Formally,

$$\text{NAS}[t] = \sum_{i=1}^{N} A(s_i[t]) \tag{4.2}$$

where $A(s_i[t])$ is the step function

$$A(s_i[t]) = \begin{cases} 1 & \text{if } s_i[t] > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{4.3}$$

(3) Weighted sum of sensor values (WSSV)

While a person will typically place their feet on the floor when sitting on the beds edge, this may not always be the case. In cases where a subject sits with legs and feet on the mattress, the sum of sensor values should be approximately equal for sitting and lying positions. To address this potential problem, the third proposed feature is a weighted sum of sensor values. The weights for each sensor are chosen based on the longitudinal sensor location; sensors located at the head or foot of the mattress receive a higher weight than those located at the middle. This weighting scheme is based on the assumption that weight distribution in lying positions will typically span all three longitudinal divisions of the mattress (i.e. head, middle, and foot), and when subjects sit at the edge of the bed, they will often sit approximately in the middle of the bed.

Mathematically the weighted sum of sensor values is defined as

$$\text{WSSV}[t] = \sum_{i=1}^{N} w(x_i) s_i[t] \tag{4.4}$$

where

$$w(x_i) = \begin{cases} 1 & \text{if } \frac{L}{3} < x_i < \frac{2L}{3} \\ \lambda & \text{otherwise,} \end{cases} \tag{4.5}$$

L is the longitudinal length of the pressure sensitive mat, and $\lambda > 1$ is a constant weight.

(4) Weighted number of active sensors (WNAS)

The fourth proposed feature is the weighted number of active sensors. As with the weighted sum of sensor values, the weights are based on longitudinal sensor location. The rationale for using this weighting scheme along with the number of active sensors is to exaggerate any relative difference between sitting and lying postures.

Mathematically this feature is defined as

$$\text{WNAS}[t] = \sum_{i=1}^{N} w(x_i) A(s_i[t]) \tag{4.6}$$

where function w is defined in equation 4.5 and function A is defined in equation 4.3.

(5) Longitudinal center of pressure (LonCP) and

(6) Lateral center of pressure (LatCP)

A simple scenario helps illustrate the justification for the next two proposed features. Consider a person seated at the edge of the bed, about to lay down. The motion of laying down will often constitute turning the body about the area of sitting while simultaneously reclining. With such a motion, the center of exerted pressure will ascend to the head of the bed as the person's torso reclines into the lying position. Thus, the center of pressure when seated will differ from the center of pressure when lying. The fifth and sixth proposed features are the centers of pressure in both the longitudinal and lateral directions. The longitudinal direction is computed as

$$\text{LonCP}[t] = \frac{1}{\text{SSV}[t]} \sum_{i=1}^{N} x_i s_i[t] \tag{4.7}$$

and the lateral direction is computed as

$$\text{LatCP}[t] = \frac{1}{\text{SSV}[t]} \sum_{i=1}^{N} y_i s_i[t] \tag{4.8}$$

(7) Longitudinal variance (LonV) and

(8) Lateral variance (LatV)

A commonly used measure of the spread of data is variance. The seventh and eighth proposed features are the active sensor location variances. This is the variance in

sensor location for active sensors in either the longitudinal or lateral direction. In the longitudinal direction

$$\mathrm{LonV}[t] = \frac{1}{\mathrm{NAS}[t]} \sum_{i=1}^{N} A(s_i[t])(x_i - \bar{x}[t])^2 \tag{4.9}$$

where the over bar denotes a sample mean (i.e. $\bar{x}[t] = \frac{1}{NAS[t]} \sum A(s_i[t])x_i$). In the lateral direction

$$\mathrm{LatV}[t] = \frac{1}{\mathrm{NAS}[t]} \sum_{i=1}^{N} A(s_i[t])(y_i - \bar{y}[t])^2. \tag{4.10}$$

4.1.2 Classifiers and hyper-parameter selection

Support vector machine (SVM), nearest neighbour (NN), and k-nearest neighbour (kNN) classifiers were chosen for this study, with all three evaluated as a candidate pattern classifier for posture recognition (see figure 4.4). The chosen implementation for SVM was the SVM module of Orange [95], a machine learning package for the Python [96] programming language. The Orange SVM module provides a Python interface to the popular libsvm C implementation [97]. Similarly, the kNN module of Orange was the chosen implementation of NN and kNN.

As discussed in section 2.1.1, using an SVM classifier requires choosing a particular kernel function. A radial basis function (RBF) kernel was chosen in this work because it can provide a nonlinear mapping from features to classes, has fewer hyper-parameters than other kernels, and is generally recommended as a first kernel to try on new problems [33]. When using an RBF kernel the SVM has two hyper-parameters which must be specified before training. The C hyper-parameter is a misclassification penalty which appears in the specification of the SVM optimization problem expressed in equation 2.1. The γ hyper-parameter appearing in equation 2.5 determines the spatial extent of the RBF in feature space and controls smoothness of the

discriminant surface.

For the SVM classifier, selection of hyper-parameters was performed using a grid-search with n-fold cross validation on the training data, as recommended in [33]. Before training the SVM on a training set, all combinations of hyper-parameter values $C \in \{2^{-5}, 2^{-4}, \ldots, 2^{15}\}$ and $\gamma \in \{2^{-15}, 2^{-14}, \ldots, 2^{3}\}$ were evaluated using n-fold cross validation on the training set and the best performing combination of C and γ were selected. The SVM was then trained using the selected values for the hyper-parameters.

The hyper-parameter for a kNN classifier is the value of k. The kNN classifier also used a hyper-parameter selection procedure; values of $k \in \{3, 5, \ldots, 29\}$ were tried in succession, and the value of k with best estimated classification accuracy estimated via n-fold cross validation was chosen. Odd numbers of k are used in binary classification tasks to avoid the possibility of a tie between classes [30].

4.1.3 Feature preprocessing

SVMs often have better performance when features are scaled to the range [-1,1] because this prevents larger scale features from dominating smaller scale features and helps avoid numerical difficulties during optimization [33]. Similarly, nearest neighbour classifiers are also sensitive to feature scale [98]. Taking these properties into consideration, the Orange SVM and kNN modules implement a simple linear scaling of each feature so that feature values fall in the range [-1,1] prior to training.

4.2 Recognition performance evaluation

Figure 4.6 depicts an overview of the basic performance evaluation methodology used to evaluate the static activity and lying posture recognizers. Pressure data was manually segmented into posture specific segments. Each posture specific segment was

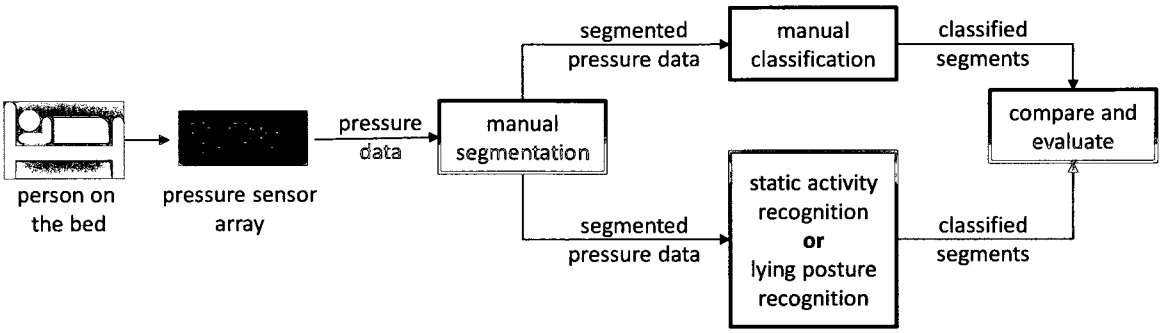

Figure 4.6: Static activity and lying posture recognition performance evaluation methodology.

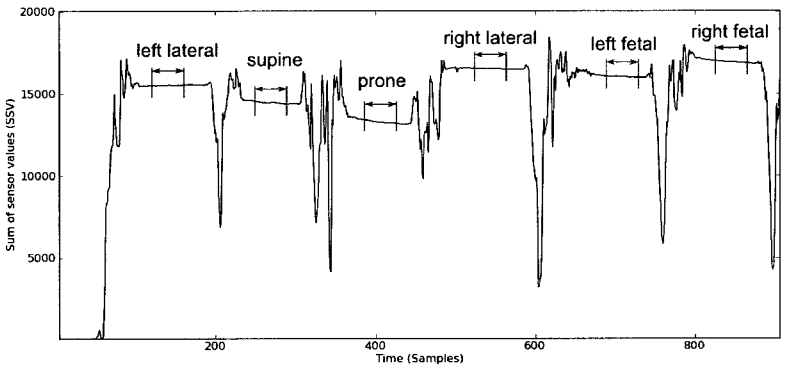

Figure 4.7: Example of manual segmentation and classification of pressure data into correct postures.

manually classified based on the known posture performance sequence. Figure 4.7 shows an example of pressure data manually segmented with the correct posture classifications. The manually segmented pressure data was input into either the static activity or lying posture recognizer, and the results from static activity or lying posture recognition were compared to the manually classified postures. The subject-dependent and subject-independent static activity data sets were used to evaluate the static activity recognizer and the lying posture data set was used to evaluate the lying posture recognizer.

4.2.1 Feature comparison and selection

To compare candidate features on an individual basis, the candidate pattern classifiers were tested and compared using each candidate feature individually. With the subject-dependent static activity data, individual feature performance was evaluated using 10-fold cross validation to estimate static activity recognition accuracy. With the subject-independent static activity data, feature performance was evaluated using leave-one-subject-out validation to estimate static activity recognition accuracy. Sensitivity and specificity were also determined for static activity recognition. With the lying posture data, feature performance was evaluated using 10-fold cross validation to estimate lying posture recognition accuracy.

As with feature comparison, feature selection was performed using subject-dependent static activity, subject-independent static activity, and lying posture data sets. Common methods for feature selection include forward selection or backward elimination [99]; however, both the total number of training examples and the maximum number of features were relatively few, so an exhaustive feature subset search was feasible, eliminating the need for a heuristic search method. To determine the best combination of features, every possible combination of features was evaluated using estimated classification accuracy as the performance metric for comparison. Classification accuracy was estimated via 10-fold cross validation in the subject-dependent case, leave-one-subject-out validation in the subject-independent case, and 10-fold cross validation in the lying posture case. This method of feature selection is commonly called a wrapper method [100].

The phrase **best feature subset of cardinality** j is here used to refer to a subset of the candidate features which had the best performance when compared with all other candidate feature subsets of cardinality j.

4.2.2 Performance assessment

The performance of the static activity recognizer and lying posture recognizer were assessed relative to previously published results in static activity recognition and lying posture recognition. Comparison to these previously published results is discussed in sections 4.3.2 and 4.3.5. The posture recognizers were also compared to two naive classifiers: the random classifier and the majority classifier.

4.3 Results and discussion

This section discusses the results of feature comparison and selection for both static activity and lying posture recognition and compares the recognition performance to previously published results.

4.3.1 Feature comparison and selection for static activity recognition

The results of feature comparison and selection for static activity recognition are presented in two sections to distinguish the subject-dependent analysis from the subject-independent analysis.

4.3.1.1 Subject-dependent analysis

Figure 4.8 depicts the results of the subject-dependent (SD) feature comparison. The figure suggests that, for most features, there was little difference in performance across classifiers. With the WSSV and LonV features, SVM performance was weaker than with either NN or kNN. Of the candidate classifiers, the SVM classifier appeared to be most sensitive to the choice of feature. This is reflected in the standard deviations of the mean classification accuracy across all features, which was 16.2% for SVM, 6.5% for NN, and 5.9% for kNN (see table 4.2).

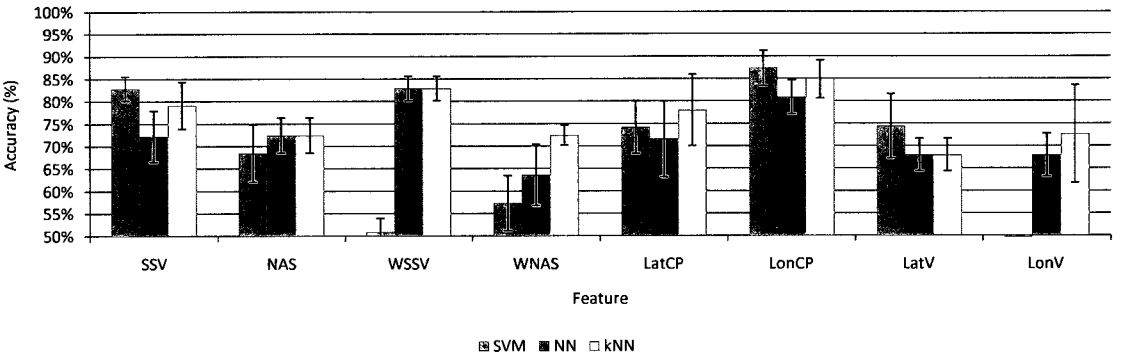

(a) Classification accuracy using individual features to recognize sitting and lying activities.

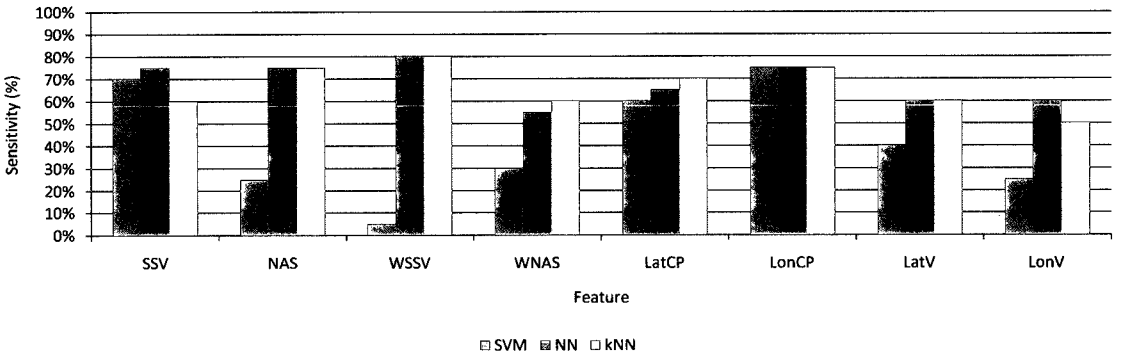

(b) Sensitivity of individual features when recognizing lying activities.

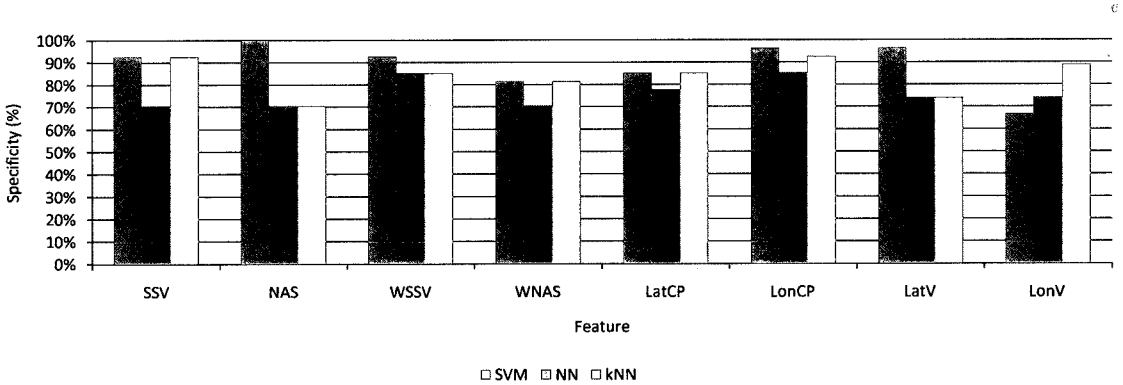

(c) Specificity of individual features when recognizing lying activities.

Figure 4.8: Results using individual features to recognize sitting and lying activities in the subject-dependent static activity data.

With all features, the NN and kNN classifiers had a classification accuracy exceeding both of the naive classifiers. The best overall estimated classification accuracy was achieved by SVM using the LonCP feature. The kNN classifier also achieved maximum performance using LonCP.

The best sensitivity with respect to lying positions was achieved using the WSSV feature with the NN and kNN classifiers. This was followed closely by the LonCP feature using any of SVM, NN, or kNN. Specificity with respect to lying positions was greatest using the NAS feature with the SVM classifier, followed by LonCP with SVM.

	SD Mean Accuracy (%)	SD Sensitivity (%)	SD Specificity (%)
SVM	16.2	24.7	10.8
NN	6.5	9.2	6.3
kNN	5.9	10.3	8.1

Table 4.2: Standard deviations of mean accuracy, sensitivity, and specificity of classifiers recognizing sitting and lying activities in the subject-dependent static activity data.

Specificity was higher than sensitivity for all features using all classifiers; the likely cause of this result was the class imbalance present in the subject-dependent data set (i.e. there was a higher proportion of sitting postures). Table 4.2 reports the variability of mean accuracy, sensitivity and specificity across the classifiers. The results show that sensitivity was more variable across features than specificity, suggesting that the variability in classification accuracy was likely the result of feature sensitivity to lying positions.

Overall, these results show that for static activity recognition reasonable classification accuracy (approximately 85%) could be achieved with a single feature if some subject-specific system training data is available.

Figure 4.9 depicts the results of subject-dependent feature selection. The plot suggests that the choice of classifier was relatively unimportant, as the curves do

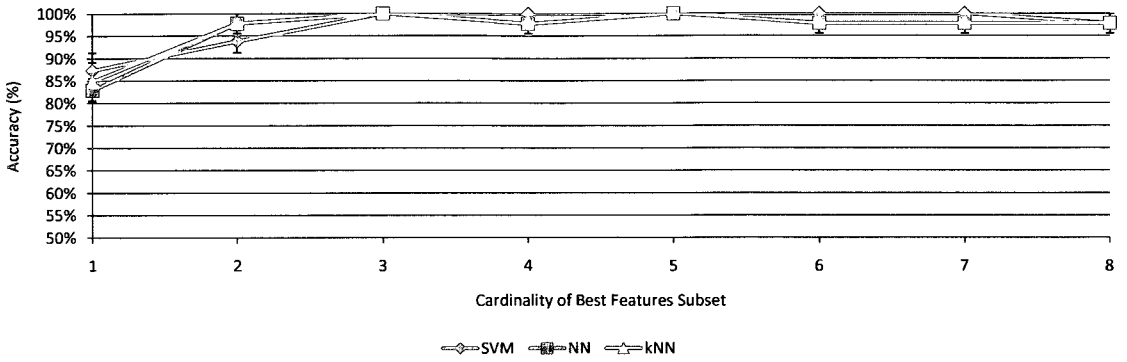

Figure 4.9: Performance of the best combinations of features for recognizing sitting and lying activities in the subject-dependent static activity data.

Cardinality	**SVM**	***k*NN**	**1NN**
1	LonCP	LonCP	WSSV
2	SSV NAS	LatCP LonV	LatCP LonV
3	WNAS LatCP LonV	LatCP LonCP LonV	LatCP LonCP LonV

Table 4.3: Best feature subsets determined using the subject-dependent static activity data. Note that all classifiers reached peak performance with 3 features (see figure 4.9), thus only best feature subsets up to cardinality 3 are displayed.

not substantially differ from each other. The naive classifiers were outperformed for each best feature subset cardinality. Using the best combination of three features, all three classifiers achieved 100% estimated classification accuracy. Using more than three features required additional computation which could not improve recognition performance.

Table 4.3 shows the features present in the best feature subsets up to cardinality 3. The NN and kNN classifiers both achieved top performance using the same three features: LonCP, LatCP, and LonV. The SVM classifier also achieved top performance with three features, including both LatCP and LonV. This suggests that, among the candidate features, the LatCP, LonCP, and LonV features were most useful for discriminating between sitting and lying. Using only LatCP and LonV, the NN and kNN classifiers achieved greater than 95% estimated classification accuracy.

Overall, these results demonstrate that an excellent static activity recognition system can be constructed using conceptually simple and computationally inexpensive features in combination with an optical pressure sensor array positioned beneath a mattress, assuming that a small amount of training data may be acquired from the person to be monitored.

4.3.1.2 Subject-independent analysis

Figure 4.10 depicts the results of subject-independent (SI) feature comparison. With most features, the SVM classifier performed best, but with only small performance gains over NN and kNN classifiers. Classification accuracy was greater with kNN than with NN for all features. Table 4.4 reports standard deviations of the mean accuracy, sensitivity with respect to lying positions, and specificity with respect to lying positions for each classifier. The results show that estimated classification accuracy was only slightly more variable across features when using SVM than the other classifiers: 13.9% standard deviation with SVM, 11.7% with NN, and 9.7% with kNN. For all classifiers, specificity with respect to lying positions was greater than sensitivity. Similar to the subject-dependent analysis, the variability of sensitivity across features was greater than the variability of specificity across features for all classifiers, indicating that the variability in estimated classification accuracy is more likely due to the features' sensitivities to the lying positions.

All features provided performance exceeding the naive classifiers. The majority of features (5 of 8) reached estimated classification accuracies exceeding 80%, which could be explained by the fixed protocol used in data collection, as participants would only lie in the supine position and would generally sit in the middle of the bed. Even so, the result demonstrates the efficacy of using a single pressure signal feature to discriminate between sitting and lying postures in a fixed posture performance protocol. The WNAS feature was the top individual performer across all performance

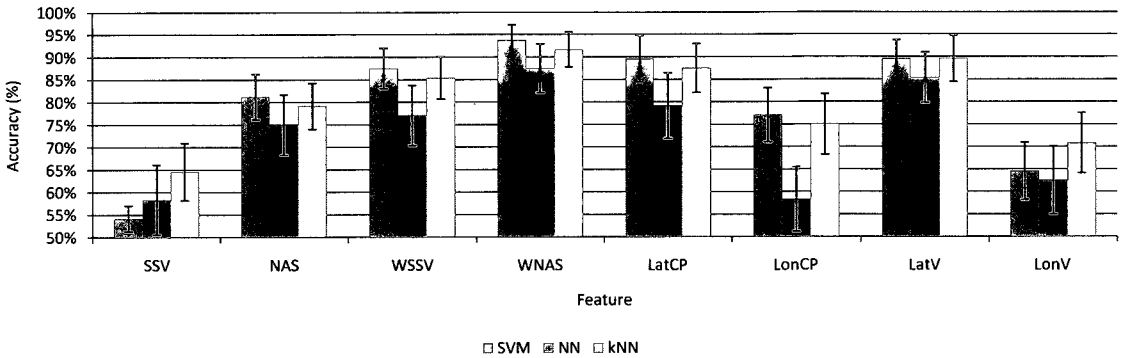

(a) Classification accuracy using individual features to recognize sitting and lying activities.

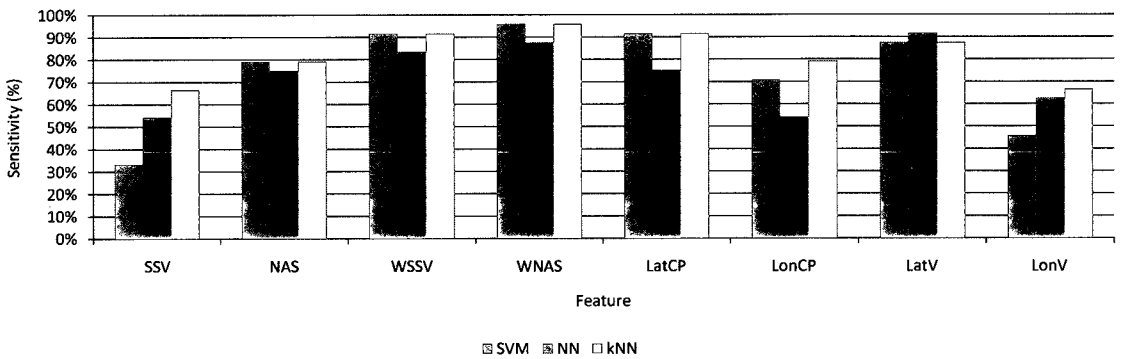

(b) Sensitivity of individual features when recognizing lying activities.

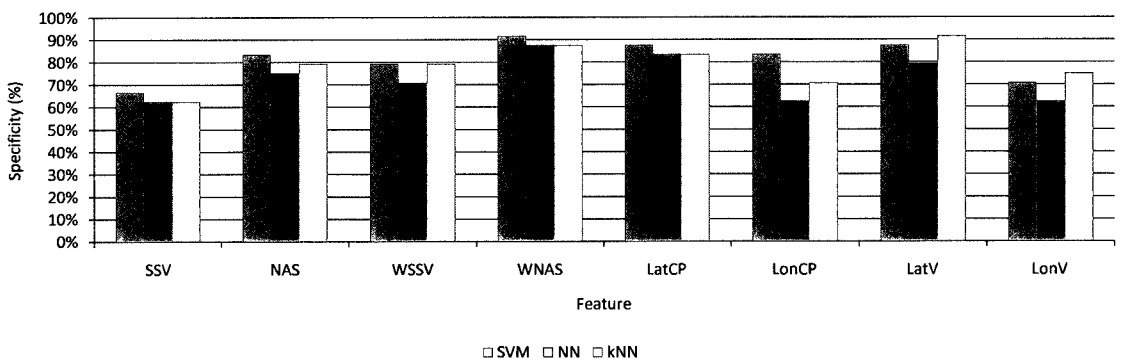

(c) Specificity of individual features when recognizing lying activities.

Figure 4.10: Results using individual features to recognize sitting and lying activities in the subject-independent static activity data.

	SI Mean Accuracy (%)	SI Sensitivity (%)	SI Specificity (%)
SVM	13.9	23.2	8.6
NN	11.7	14.6	10.0
kNN	9.7	11.3	9.3

Table 4.4: Standard deviations of mean accuracy, sensitivity, and specificity of classifiers recognizing sitting and lying activities in the subject-independent static activity data.

metrics (accuracy, sensitivity, and specificity). The weighting scheme for WSSV and WNAS was specifically designed to exaggerate the difference in the sum of active sensor values and number of active sensors, respectively, between a *typical* lying position (where the subject lies extended in the longitudinal direction of the mattress) and a *typical* sitting position (where the subject sits on the middle of the mattress longitudinally). The greater performance achieved by the WSSV and WNAS features in comparison to the SSV and NAS in the subject-independent feature comparison validates the weighting scheme design choice.

The results suggest that recognition accuracies above 90% are achievable, using only one of the candidate signal features, in data collection scenarios with a fixed posture performance protocol and no subject-specific system training.

Figure 4.11 depicts the results of subject-independent feature selection. As in the subject-dependent analysis, the choice of classifier appeared to have little importance; the performance curves for each classifier had approximately the same shape. SVM performance was consistently better than both NN and kNN, but only marginally so. Overall, classification performance significantly exceeded the naive classifiers. With five features (NAS, WSSV, LatCP, LatV, and LonV) and the SVM classifier, estimated classification accuracy reached 100%. The NN and kNN classifiers both achieved their best performance with four features.

Table 4.6 reports the features included in the best feature subsets up to cardinality 3. A review of the table shows that LatCP and WNAS were important features for

	SVM	kNN	NN
1	WNAS	WNAS	WNAS
2	WSSV LatCP	WSSV LatCP	WSSV LatCP
	WNAS LatCP	WNAS LatCP	
3	SSV WNAS LatCP	SSV WNAS LatCP	SSV WNAS LatCP
	SSV WSSV LonV	NAS WNAS LonCP	NAS WNAS LonCP
	WNAS LatCP LatV	WSSV LatCP LonCP	NAS WNAS LatCP
	SSV WSSV LatV	NAS WNAS LatCP	WSSV WNAS LatCP
	SSV WSSV LatCP	WNAS LatCP LatV	SSV WSSV LatCP
		WNAS LatV LonV	WNAS LatCP LonCP
		LatCP LonCP LonV	
		WSSV WNAS LatCP	
		SSV WSSV LatCP	
		WNAS LatCP LonV	
		WNAS LatCP LonCP	

Table 4.6: Best feature subsets determined using the subject-independent static activity data.

static activity recognition with the subject-independent data. WNAS appeared in all best feature subsets of cardinality 1, 2 out of 5 (40%) of best feature subsets of cardinality 2, and 15 out of 22 (68.2%) of best feature subsets of cardinality 3. Similarly, LatCP appeared in all best feature subsets of cardinality 2, and 17 out of 22 (77.3%) best feature subsets of cardinality 3. WSSV was also prominent in the best feature subsets, most notably in the best feature subsets of cardinality 2.

The static activity recognizer is able to accurately recognize sitting and lying postures unobtrusively using subject-independent training and a fixed posture performance protocol. These constraints – subject-independent training and fixed posture protocol – are realistic; they could manifest in rehabilitation scenarios in which no previous training data for the patient has been acquired and the patient is asked to perform particular movements such as basic activities of daily living (e.g. sit-to-lie postural transitions). For such scenarios, the results presented in this section demonstrate that accurate posture classification under these constraints is achievable using

Reference	Sensors	Best subject-dependent result (%)	Best subject-independent result (%)
Present work	**Optical bed-based pressure sensor array**	**100**	**100** (note: participants were performing a fixed protocol and only lay in the supine position)
Tapia et al. [68]	3-axis accelerometers heart-rate monitor	Precision: 98.9 Recall: 99.2	Precision: 86.1 Recall: 82.3
Baek et al. [63]	3-axis accelerometer	100	n/a
Lee et al. [64]	3-axis accelerometer	n/a	94.0-96.6
Najafi et al. [79]	2-axis accelerometer gyroscope	n/a	Sitting sensitivity: 90.2 Sitting specificity: 93.4 Lying sensitivity: 98.4 Lying specificity: 99.7
Ermes et al. [65]	3-axis accelerometers	n/a	Sensitivity: 97
Culhane et al. [66]	2-axis accelerometers	n/a	Sitting: 92 Lying: 95
Lyons et al. [67]	2-axis accelerometers	n/a	Sitting: 93 Lying: 84

Table 4.7: Comparison to performance achieved in published research for recognizing sitting and lying static activities. Where the particular performance metric is not specified it is classification accuracy.

conceptually simple and computationally inexpensive pressure signal features with an unobtrusive bed-based pressure sensor array.

4.3.2 Comparison to previously published static activity detection results

Previously published results serve as the external standard used to judge the success of the static activity recognizer described in this work. The studies used for comparison employ a range of performance metrics and evaluation techniques but typically report one or more of the related performance metrics including classification accuracy, precision, recall, sensitivity, and specificity.

Table 4.7 summarizes the types of sensors used, best subject-dependent results, and best subject-independent results reported in the recent literature on activity and posture recognition using wearable sensors. While most of these studies investigated both static and dynamic activities, table 4.7 only reports on recognition of sitting and lying activities.

With subject-dependent analysis the static activity recognizer achieved 100% classification accuracy and matched the best subject-dependent result achieved using wearable sensors as reported in [63]. This result means that a bed-based pressure sensor array could be used in place of wearable sensors for reliably recognizing if a person is sitting or lying in bed when the system has been trained to recognize this particular person.

With subject-independent analysis the static activity recognizer again achieved 100% classification accuracy, outperforming the published techniques using wearable sensors. While this result is encouraging, recall that the lying activities used in the subject-independent analysis were fixed to supine lying postures; as such, the subject-independent results should not be overstated. However, having achieved 100% classification accuracy, it is reasonable to claim that the subject-independent performance was competitive with the performances reported in the wearable sensor literature. Even if no prior training data is available for a person, a bed-based optical pressure sensor array could replace wearable sensors to reliably recognize sitting and lying activities while the person is in bed using the techniques described.

4.3.3 Feature comparison and selection for lying posture recognition

Figure 4.12 depicts the estimated classification accuracy using individual features for the lying posture recognition task. The figure indicates that for most features, there was little difference in performance across classifiers. For all classifiers, the performance of each candidate feature exceeded the performance of both naïve classifiers. The overall best performance of 50.0 $\pm$ 6.1% estimated accuracy was achieved using the kNN classifier with $k = 15$ and the LonCP feature. Using the NN classifier, estimated performance reached a peak of 38.6 $\pm$ 5.2% with the LonV feature. The SVM classifier achieved 44.3 $\pm$ 2.6% estimated accuracy using the WSSV feature.

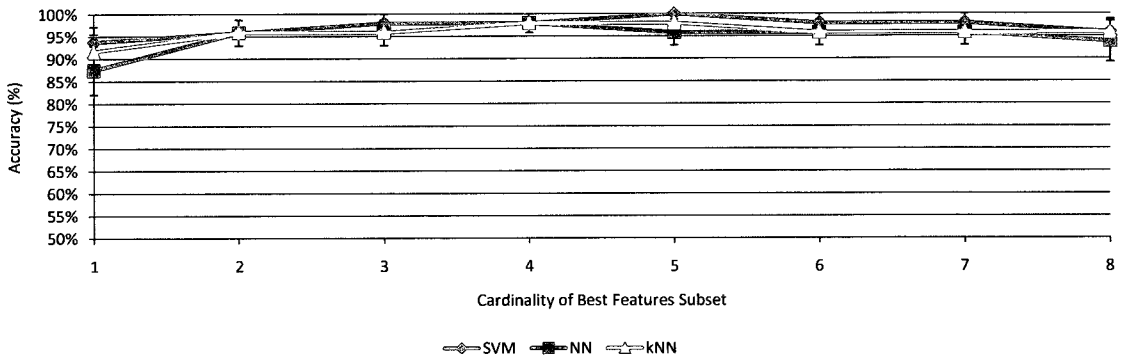

Figure 4.11: Performance of the best combinations of features for recognizing sitting and lying activities in the subject-independent static activity data.

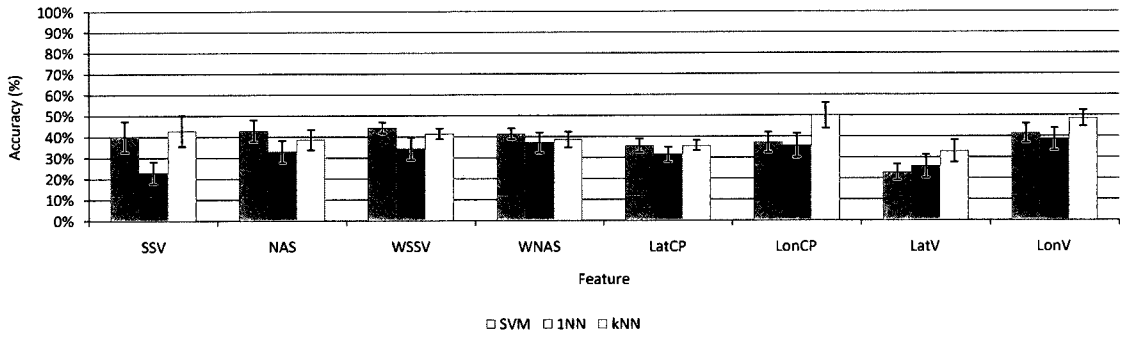

Figure 4.12: Classification accuracy using individual features to recognize lying postures.

Actual Class \ Predicted Class	LL	RL	Su	Pr	LF	RF	Si	Recall
LL	**5**	1	0	0	0	4	0	0.500
RL	0	**8**	0	0	0	2	0	0.800
Su	0	2	**0**	0	4	0	4	0.000
Pr	0	3	0	**4**	0	2	1	0.400
LF	0	0	0	0	**8**	2	0	0.800
RF	0	4	0	0	4	**2**	0	0.200
Si	1	2	1	0	1	1	**4**	0.400
Precision	0.833	0.400	0.000	1.000	0.471	0.154	0.444	

(a) SVM with the WSSV feature

Actual Class \ Predicted Class	LL	RL	Su	Pr	LF	RF	Si	Recall
LL	**3**	0	4	1	0	0	2	0.300
RL	0	**3**	4	1	0	2	0	0.300
Su	3	4	**0**	2	1	0	0	0.000
Pr	1	3	0	**4**	2	0	0	0.400
LF	0	1	1	2	**4**	2	0	0.400
RF	0	2	0	0	2	**6**	0	0.600
Si	3	0	0	0	0	0	**7**	0.700
Precision	0.300	0.231	0.000	0.400	0.444	0.600	0.778	

(b) NN with the LonV feature

Actual Class \ Predicted Class	LL	RL	Su	Pr	LF	RF	Si	Recall
LL	**4**	0	3	2	0	0	1	0.400
RL	1	**7**	2	0	0	0	0	0.700
Su	1	3	**4**	1	0	0	1	0.400
Pr	2	2	2	**3**	0	0	1	0.300
LF	0	0	0	0	**5**	5	0	0.500
RF	0	0	0	0	4	**6**	0	0.600
Si	0	2	0	2	0	0	**6**	0.600
Precision	0.500	0.500	0.364	0.375	0.556	0.545	0.667	

(c) kNN with the LonCP feature

Table 4.8: Resulting confusion matrices for SVM, NN, and kNN using the best single feature to recognize lying postures.

Table 4.8 shows the resulting confusion matrices using the individually best performing feature for each candidate classifier. The SVM classifier using the WSSV feature (table 4.8a) had poor recognition performance with the supine, right fetal, and sitting up positions. The supine position was always misclassified as one of right lateral (2 misclassifications), left fetal (4), or sitting up (4); the right fetal position was only correctly identified twice, and was misclassified as right lateral (4) or left fetal (4); the sitting up position was misclassified at least once as each of the other positions, with the exception of the prone position. The SVM classifier was most precise at identifying the prone position, not once was a position misclassified as prone. The SVM classifier achieved the best recall rates on the right lateral and left fetal positions. Overall, SVM with WSSV was biased towards the right lateral and left fetal positions, classifying approximately 29% of the positions as right lateral, 24% as left fetal, and 19% as right fetal which accounts for almost three quarters of the predicted positions. The majority of misclassifications were due to this bias in favor of the right lateral and left fetal positions.

The NN classifier using LonV (table 4.8b) had poor recall rates for all positions except right fetal and sitting up. This result supports the intuition used in choosing

the LonV feature for discriminating between sitting and lying positions: the spread of pressure in the longitudinal direction should be less with sitting positions than lying positions. As with SVM, the supine position was never correctly classified. In contrast to SVM, the right lateral and left fetal positions were often misclassified. The NN classifier was less biased to particular positions than SVM, predicting each class 9-13% of the time.

The kNN classifier (table 4.8c) using LonCP and $k = 15$ had poor recall rates for the left lateral, supine, and prone positions. The highest recall rate occurred with the right lateral position and the highest precision with the sitting up position. The supine and prone positions were most often confused with each other or classified as a lateral position. Similarly, when the lateral positions were misclassified they were either confused with each other or classified as either supine or prone. When the right and left fetal positions were misclassified, they were in all cases confused with each other.

Figure 4.13 depicts the results of the feature selection process for the lying position recognition task. As in the feature comparison results, the plot indicates that the choice of classifier had only a small effect on the lying posture recognizer's classification accuracy. The combination of features, however, had a large impact on accuracy. With the best combination of four candidate features, the SVM classifier achieved 91.4 $\pm$ 2.3% estimated classification accuracy. The NN and kNN classifiers fell just below this performance, both achieving a peak 88.6 $\pm$ 2.9% estimated classification accuracy with four features. The SVM classifier reached its peak performance of 92.9 $\pm$ 2.4% with all eight candidate features, although only marginally better than the performance it achieved with four features. The plot demonstrates that integrating additional features beyond a combination of four provided diminishing returns with respect to the additional computation required.

Table 4.9 shows the best feature subsets for each classifier up to a cardinality

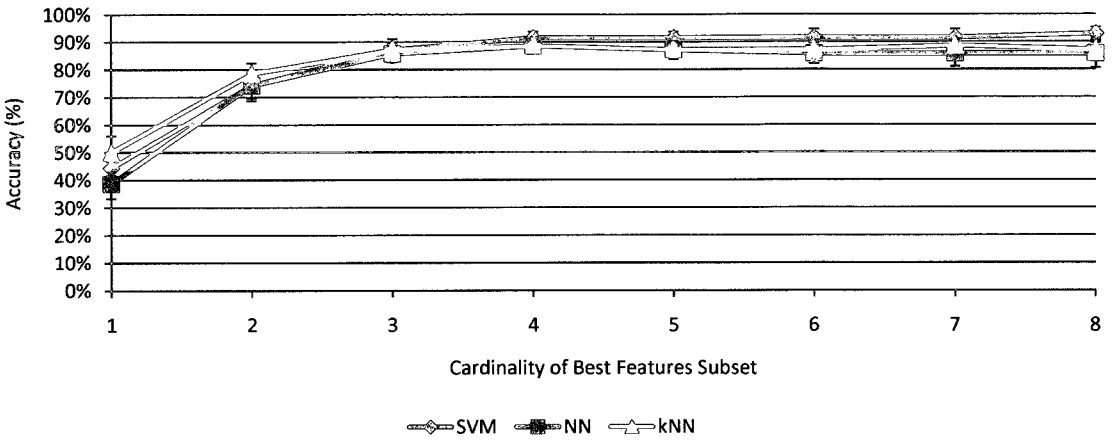

Figure 4.13: Performance of the best combinations of features for recognizing lying postures.

Cardinality	SVM	NN	kNN
1	WSSV	LonV	LonCP
2	WSSV, LonCP	WSSV, LonCP	WSSV, LonCP
3	NAS, WSSV, LonCP	WSSV, WNAS, LonCP	NAS, WSSV, LonCP
4	SSV, WNAS, LonCP, LonV	SSV, NAS, WNAS, LonCP	SSV, NAS, WNAS, LonCP

Table 4.9: Best feature subsets determined using the lying posture data.

of 4. With all three classifiers, the best performing combination of two features was WSSV and LonCP. With four features, the best feature subsets included SSV, LonCP, WNAS, and one of NAS or LonV. These features cover three basic physical properties of the interface between the person and the pressure sensors: the total exerted pressure, the location of the concentration of pressure, and the pressure distribution.

The resulting confusion matrices for SVM when using the best combinations of two and four features are shown in table 4.10. With the best combination of two features (table 4.10a), the distribution of the SVM classifier's predictions more clearly reflects the true class distribution than with the best single feature. Precision was most improved with the right lateral, supine, and right fetal positions, and recall was most improved with the supine and right fetal positions. With the best combination of four features (table 4.10b), precision was most improved with the prone and sitting up positions, and recall was best improved with the prone, right fetal, and sitting

Actual Class \ Predicted Class	LL	RL	Su	Pr	LF	RF	Si	Recall
LL	7 (+2)	2 (+1)	0	1 (+1)	0	0 (-4)	0	0.700
RL	1 (+1)	9 (+1)	0	0	0	0 (-2)	0	0.900
Su	1 (+1)	0 (-2)	9 (+9)	0	0 (-4)	0	0 (-4)	0.900
Pr	1 (+1)	0 (-3)	2 (+2)	5 (+1)	0	0 (-2)	2 (+1)	0.500
LF	0	0	0	0	8	2	0	0.800
RF	0	0 (-4)	0	0	3 (-1)	7 (+5)	0	0.700
Si	0 (-1)	0 (-2)	1	2 (+2)	0 (-1)	0 (-1)	7 (+3)	0.700
Precision	0.700	0.818	0.750	0.625	0.727	0.778	0.778	

(a) SVM with best two features

Actual Class \ Predicted Class	LL	RL	Su	Pr	LF	RF	Si	Recall
LL	9 (+2)	0 (-2)	1 (-1)	0 (-1)	0	0	0	0.900
RL	0 (-1)	10 (+1)	0	0	0	0	0	1.000
Su	1	0	8 (-1)	1 (+1)	0	0	0	0.800
Pr	0 (-1)	0	1 (-1)	9 (+4)	0	0	0 (-2)	0.900
LF	0	0	0	0	8 (0)	2	0	0.800
RF	0	0	0	0	0 (-3)	10 (+3)	0	1.000
Si	0	0	0 (-1)	0 (-2)	0	0	10 (+3)	1.000
Precision	0.900	1.000	0.800	0.900	1.000	0.833	1.000	

(b) SVM with best four features

Table 4.10: Resulting confusion matrices for SVM using the best combination of two and four features to recognize lying postures. Changes are shown in parentheses.

Actual Class \ Predicted Class	LL	RL	Su	Pr	LF	RF	Si	Recall
LL	9 (+6)	0	0 (-4)	1	0	0	0 (-2)	0.900
RL	1 (-1)	9 (+6)	0 (-4)	0 (-1)	0	0 (-2)	0	0.900
Su	0 (-3)	0 (-4)	9 (+9)	0 (-2)	0 (-1)	0	1 (-1)	0.900
Pr	1	0 (-3)	2 (-2)	4 (0)	0 (-2)	0	3 (+3)	0.400
LF	0	0 (-1)	0 (-1)	0 (-2)	6 (+2)	4 (-2)	0	0.600
RF	0	0 (-2)	0	0	3 (-1)	7 (+1)	0	0.700
Si	0 (-3)	0	0	2 (+2)	0	0	8 (+1)	0.800
Precision	0.818	1.000	0.818	0.571	0.667	0.636	0.667	

(a) NN with best two features

Actual Class \ Predicted Class	LL	RL	Su	Pr	LF	RF	Si	Recall
LL	9 (0)	0	1 (-1)	0 (-1)	0	0	0	0.900
RL	0 (-1)	9 (0)	0	0	0	1 (+1)	0	0.900
Su	0	0	8 (+1)	2 (-2)	0	0	0 (-1)	0.800
Pr	0 (-1)	0	2	8 (+4)	0	0	0 (-3)	0.800
LF	0	0	0	0	10 (+4)	0 (-4)	0	1.000
RF	0	0	0	0	2 (-1)	8 (+1)	0	0.800
Si	0	0	0	0 (-2)	0	0	10 (+2)	1.000
Precision	1.000	1.000	0.727	0.800	0.833	0.889	1.000	

(b) NN with best four features

Table 4.11: Resulting confusion matrices for NN using the best combination of two and four features to recognize lying postures. Changes are shown in parentheses.

up positions. The remaining misclassifications were the result of the left and right fetal positions being confused, the supine and prone positions being confused, and one supine position misclassified as left lateral.

The resulting confusion matrices for NN when using the best combinations of two and four features are shown in table 4.11. With the best combination of two features (table 4.11a), the best improvement in recall rate occurred with the left lateral, right lateral, and supine positions. The most improvement in precision was achieved on the right lateral, supine, and left lateral positions. Using the best combination of four features (table 4.11b), the greatest improvement in recall was achieved with the prone and left fetal positions, and the greatest improvement in precision was achieved with the sitting up, right fetal, and prone positions. The remaining misclassifications fell into four categories: confusion between supine and prone, between left lateral and supine, between right lateral and right fetal, and between right fetal and left fetal.

The resulting confusion matrices for kNN when using the best combinations of two and four features is shown in table 4.12. With the best combination of two features

Actual Class \ Predicted Class	LL	RL	Su	Pr	LF	RF	Si	Recall
LL	9 (+5)	0	0 (-3)	1 (-1)	0	0	0 (-1)	0.900
RL	1	9 (+2)	0 (-2)	0	0	0	0	0.900
Su	1	1 (-2)	7 (+3)	1	0	0	0 (-1)	0.700
Pr	1 (-1)	0 (-2)	2	7 (+4)	0	0	0 (-1)	0.700
LF	0	0	0	0	8 (+3)	2 (-3)	0	0.800
RF	0	0	0	0	3 (-1)	7 (+1)	0	0.700
Si	0	0 (-2)	1 (+1)	1 (-1)	0	0	8 (+2)	0.800
Precision	0.750	0.900	0.700	0.700	0.727	0.778	1.000	

(a) kNN with best two features

Actual Class \ Predicted Class	LL	RL	Su	Pr	LF	RF	Si	Recall
LL	9 (0)	0	1 (+1)	0 (-1)	0	0	0	0.900
RL	0 (-1)	9 (0)	0	0	0	1 (−1)	0	0.900
Su	0 (-1)	0 (-1)	8 (+1)	2 (−1)	0	0	0	0.800
Pr	0 (-1)	0	2	8 (+1)	0	0	0	0.800
LF	0	0	0	0	10 (+2)	0 (-2)	0	1.000
RF	0	0	0	0	2 (-1)	8 (+1)	0	0.800
Si	0	0	0 (-1)	0 (-1)	0	0	10 (+2)	1.000
Precision	1.000	1.000	0.727	0.800	0.833	0.889	1.000	

(b) kNN with best four features

Table 4.12: Resulting confusion matrices for kNN using the best combination of two and four features to recognize lying postures. Changes are shown in parentheses.

(table 4.12a), the best improvement in recall rate occurred with the left lateral and prone positions, and the most improvement in precision was achieved on the right lateral and sitting up positions. With the best combination of four features (table 4.12b), recall was most improved with the sitting up and left fetal positions. The best improvement in precision was found in the left lateral position. The remaining misclassifications were confusions between supine and prone, between left lateral and supine, between right lateral and right fetal, and between right fetal and left fetal.

Comparing the results from all three classifiers using the best combination of four features, the general causes of misclassification were (in order of most to least common): confusion between supine and prone, confusion between right fetal and left fetal, confusion between lateral positions and supine, and confusion between lateral positions and fetal positions.

4.3.3.1 Feature comparison and selection for lying posture recognition with a reduced set of postures

Depending on the intended application of a lying posture recognition system, the capability to distinguish between right and left lateral decubitus or right and left fetal may not necessarily be required. This section evaluates the performance of the lying posture recognizer with a reduced class set where left/right lateral positions and left/right fetal postures are simplified into lateral and fetal, respectively. In other words, no distinction is made for the side of the body on which the person rests.

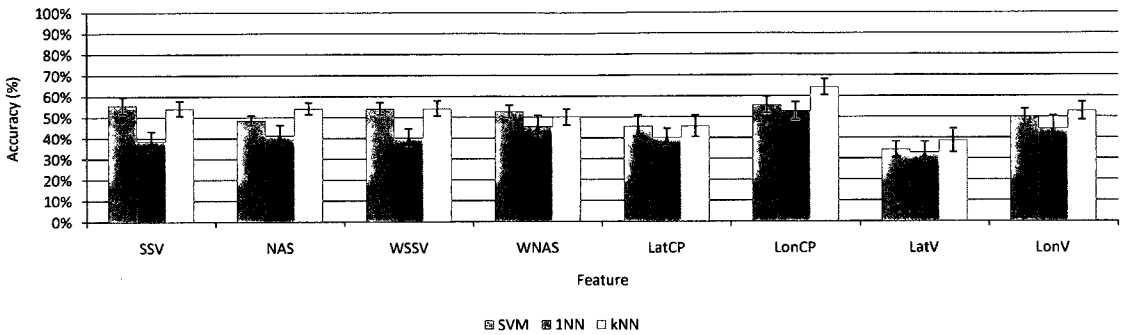

Figure 4.14: Classification accuracy using individual features to recognize lying postures with a reduced set of postures.

Figure 4.14 depicts the estimated classification accuracy of each candidate feature and classifier. Again, the plot suggests that the choice of classifier has little influence on accuracy. With every feature, each classifier outperformed both naive classifiers. The best performance of 64.3 +/- 3.8% estimated classification accuracy was attained with the kNN classifier with $k = 25$ using the LonCP feature. The SVM classifier had best performance using either the LonCP or SSV feature, in both cases achieving 55.7 +/- 4.0% estimated classification accuracy. NN reached 52.9 +/- 4.3% estimated classification accuracy, also with the LonCP feature.

Table 4.13 shows the resulting confusion matrices for each classifier using the best individual feature (in this case, LonCP for all three classifiers). The SVM classifier (table 4.13a) exhibited a strong bias towards the lateral and fetal positions, each of which was represented twice as often as the other classes in the data set. The lateral position was predicted for 42 of 70 instances (60%), and 25 of these predictions were misclassifications. Recall of the supine and prone classes was very poor; the supine position was never predicted by the classifier.

The NN classifier (table 4.13b) did not exhibit the same class bias as the SVM classifier, with predictions being more accurately divided amongst the five classes. Notably, the NN classifier with LonCP achieved perfect recall and precision with the fetal position. As with SVM, recall and precision with the supine and prone positions

		Predicted Class					
		Si	**L**	**F**	**Su**	**Pr**	**Recall**
Actual Class	**Si**	**5**	3	0	0	2	0.50
	L	1	**18**	0	0	1	0.90
	F	0	5	**15**	0	0	0.75
	Su	1	8	0	**0**	1	0.00
	Pr	1	8	0	0	**1**	0.10
	Precision	0.63	0.43	1.00	0.00	0.20	

(a) SVM with LonCP

		Predicted Class					
		Si	**L**	**F**	**Su**	**Pr**	**Recall**
Actual Class	**Si**	**6**	2	0	1	1	0.60
	L	2	**7**	0	6	5	0.35
	F	0	0	**20**	0	0	1.00
	Su	1	3	0	**3**	3	0.30
	Pr	1	6	0	2	**1**	0.10
	Precision	0.60	0.39	1.00	0.25	0.10	

(b) NN with LonCP

		Predicted Class					
		Si	**L**	**F**	**Su**	**Pr**	**Recall**
Actual Class	**Si**	**6**	2	0	0	2	0.60
	L	1	**15**	0	2	2	0.75
	F	0	1	**19**	0	0	0.95
	Su	1	8	0	**1**	0	0.10
	Pr	0	6	0	0	**4**	0.40
	Precision	0.75	0.47	1.00	0.33	0.50	

(c) kNN with LonCP

Table 4.13: Resulting confusion matrices for SVM, NN, and kNN using the best single feature to recognize lying postures with a reduced set of postures.

was poor. Unlike with SVM, recall of lateral positions was also poor.

As with NN, the kNN classifier (table 4.13c) was an excellent fetal position recognizer, with no false fetal classifications at all and just one missed fetal position out of twenty. The kNN classifier had much better recall with lateral positions than NN, but overall it exhibited a similar bias towards the lateral position as with SVM, predicting the lateral position 45.7% of the time. The kNN classifier fared no better than SVM or NN with the supine and prone positions; again, both recall and precision with these positions was poor.

Figure 4.15 depicts the feature selection results using the reduced class set. The top performance of 97.1 +/- 1.9% estimated classification accuracy was achieved by SVM using the best combination of four features. Performance began to decline with greater than five features for all three classifiers. With greater than one feature, the performance differences between classifiers were small.

Table 4.14 shows the best feature subsets up to cardinality 5 for all three classifiers, and tables 4.15, 4.16, and 4.17 show resulting confusion matrices using the best two and best four features for SVM, NN, and kNN, respectively.

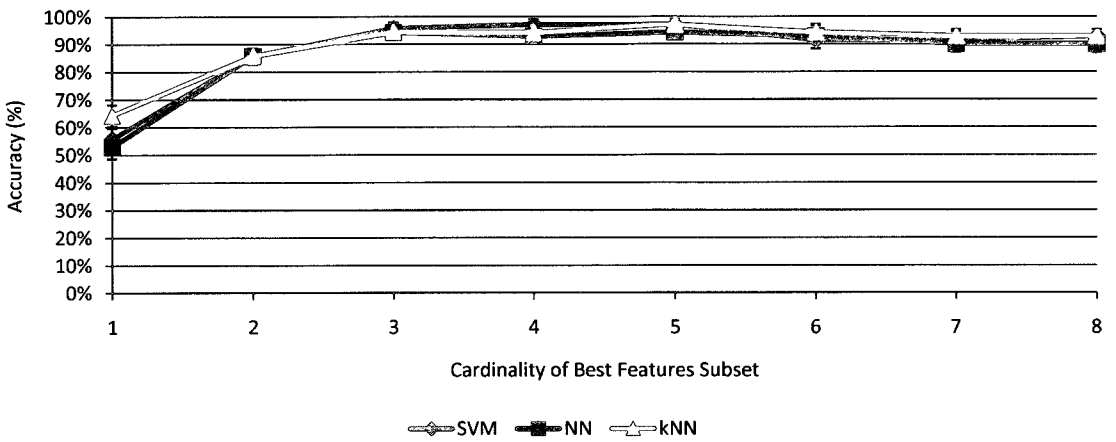

Figure 4.15: Performance of the best combinations of features for recognizing lying postures with a reduced set of postures

Cardinality	SVM	NN	kNN
1	LonCP	LonCP	LonCP
	SSV		
2	SSV, LonCP	WSSV, LonCP	WSSV, LonCP
3	SSV, NAS, LonCP	SSV, NAS, LonCP	SSV, NAS, LonCP
			SSV, LonCP, LatV
4	SSV, WSSV, LatCP, LonCP	SSV, LatCP, LonCP, LonV	SSV, NAS, WSSV, LonCP
			SSV, NAS, LatCP, LonCP
			SSV, WSSV, LonCP, LonV
5	NAS, WSSV, LatCP, LonCP, LatV	SSV, NAS, LatCP, LonCP, LatV	SSV, NAS, LatCP, LonCP, LatV

Table 4.14: Best feature subsets determined using the lying posture data with a reduced set of postures.

Actual Class \ Predicted Class	Si	L	F	Su	Pr	Recall
Si	6 (+1)	0 (-3)	0	0	4 (+2)	0.60
L	0 (-1)	19 (+1)	0	1 (+1)	0 (-1)	0.95
F	0	0 (-5)	20 (+5)	0	0	1.00
Su	0 (-1)	1 (-7)	0	8 (+8)	1	0.80
Pr	0 (-1)	0 (-8)	0	3 (+3)	7 (+6)	0.70
Precision	1.00	0.95	1.00	0.67	0.58	

(a) SVM with best two features

Actual Class \ Predicted Class	Si	L	F	Su	Pr	Recall
Si	10 (+4)	0	0	0	0 (-4)	1.00
L	0	20 (+1)	0	0 (-1)	0	1.00
F	0	0	20	0	0	1.00
Su	0	0 (-1)	0	9 (+1)	1	0.90
Pr	0	0	0	1 (-2)	9 (+2)	0.90
Precision	1.00	1.00	1.00	0.90	0.90	

(b) SVM with best four features

Table 4.15: Resulting confusion matrices for SVM using the best combination of two and four features to recognize lying postures with a reduced set of postures. Changes are shown in parentheses.

The best single feature for SVM was LonCP, and the best combination of two features was SSV and LonCP. With the addition of SSV, the SVM class prediction distribution was much closer to the true class distribution than with LonCP alone. Notably, the SVM classifier achieved perfect recall and precision with the fetal position. Recall of the supine and prone positions was significantly improved, however precision with supine and prone was well below the precision achieved with the other three classes. The sitting up position had the lowest recall, but perfect precision. The best combination of four features with SVM included SSV, WSSV, LatCP, and LonCP. With these features, perfect precision and recall was achieved with the sitting up, lateral, and fetal positions. The only misclassifications occurred with the supine and prone positions; one prone position was misclassified as supine and one supine position was misclassified as prone. It is clear that the most difficult task for the lying posture recognizer was to discriminate between supine and prone positions.

With the NN classifier, the best combination of two features was WSSV and LonCP. As table 4.16a shows, both recall and precision with lateral positions was significantly improved with the addition of WSSV, and the NN classifier remained a perfect fetal recognizer. Recall and precision with supine positions was also significantly improved, and slightly improved with prone positions. Performance of the NN classifier with WSSV and LonCP reached an estimated classification accuracy of 85.7 +/- 3.0%. The NN classifier reached a peak performance of 94.3 +/- 2.3% with three

Actual Class \ Predicted Class	Si	L	F	Su	Pr	Recall
Si	8 (+2)	0 (-2)	0	0 (-1)	2 (-1)	0.80
L	0 (-2)	19 (+12)	0	0 (-6)	1 (-4)	0.95
F	0	0	20	0	0	1.00
Su	1	0 (-3)	0	9 (+6)	0 (-3)	0.90
Pr	3 (+2)	1 (-5)	0	2	4 (+3)	0.40
Precision	0.67	0.95	1.00	0.82	0.57	

(a) NN with best two features

Actual Class \ Predicted Class	Si	L	F	Su	Pr	Recall
Si	10 (+2)	0	0	0	0 (-2)	1.00
L	0	20 (+1)	0	0	0 (-1)	1.00
F	0	0	20	0	0	1.00
Su	0 (-1)	0	0	8 (-1)	2 (+2)	0.80
Pr	0 (-3)	0 (-1)	0	2	8 (+4)	0.80
Precision	1.00	1.00	1.00	0.80	0.80	

(b) NN with best three features

Table 4.16: Resulting confusion matrices for NN using the best combination of two and three features to recognize lying postures with a reduced set of postures. Changes are shown in parentheses.

features: SSV, NAS, and LonCP. Table 4.16b demonstrates that with these features the NN classifier had perfect precision and recall with the sitting up, lateral, and fetal positions. As with the SVM classifier, the most difficult task was discriminating between supine and prone; all four misclassifications were confusions between these positions.

The kNN classifier's best performance using a combination of two features also occurred with WSSV and LonCP. Using these two features, the kNN classifier achieved an 85.7 +/- 2.1% estimated classification accuracy. When WSSV was combined with LonCP, recall and precision were significantly improved with the lateral and supine positions and reached 100% with the fetal position. Recall and precision were poorest with the supine and prone positions. The kNN classifier had its best performance using five features (see table 4.14), achieving 97.1 +/- 1.9% estimated classification accuracy. With the best combination of five features, the kNN classifier only misclassified two positions: one fetal position was misclassified as a lateral position, and a supine position was misclassified as prone.

Several conclusions may be drawn from the results using a reduced class set. The confusion matrices indicated that the fetal position was the most straightforward to discriminate from other positions. This result is encouraging, as intuition suggests that the pressure distribution generated by a subject's body in the fetal position should be quite distinct from the positions in which the body is stretched longitu-

		Predicted Class					
		Si	L	F	Su	Pr	Recall
Actual Class	Si	8 (+2)	0 (-2)	0	0	2	0.80
	L	0 (-1)	19 (+4)	0	0 (-2)	1 (-1)	0.95
	F	0	0 (-1)	20 (+1)	0	0	1.00
	Su	1	1 (-7)	0	7 (+6)	1 (+1)	0.70
	Pr	1 (-1)	1 (-5)	0	2 (-2)	6 (+2)	0.60
	Precision	0.80	0.90	1.00	0.78	0.60	

(a) kNN with best two features

		Predicted Class					
		Si	L	F	Su	Pr	Recall
Actual Class	Si	10 (+2)	0	0	0	0 (-2)	1.00
	L	0	20 (+1)	0	0	0 (-1)	1.00
	F	0	1 (−1)	19 (-1)	0	0	0.95
	Su	0 (-1)	0 (-1)	0	9 (−2)	1	0.90
	Pr	0 (-1)	0 (-1)	0	0 (-2)	10 (+4)	1.00
	Precision	1.00	0.95	1.00	1.00	0.91	

(b) kNN with best five features

Table 4.17: Resulting confusion matrices for kNN using the best combination of two and five features to recognize lying postures with a reduced set of postures. Changes are shown in parentheses.

dinally. If accurate fetal position detection was the primary practical goal for the implementation of such a system, then a conceptually simple fetal posture recognizer could be implemented using the NN classifier with the LonCP feature alone.

Distinguishing between prone and supine positions was a difficult task. Achieving perfect accuracy with either or both of the prone and supine positions likely requires either a higher resolution sensor or more sophisticated pressure analysis, perhaps by considering the time variation of pressure caused by breathing.

If maximum classification accuracy across all position types was the primary goal of the system's implementation, and a fairly large set of training data for the subject to be monitored was available, then these results suggest that the SVM classifier with the SSV, WSSV, LatCP, and LonCP features would be the best choice. However, accurate classification with a minimum amount of training is desirable, as the target population for such devices would be the infirm and elderly who may not have the strength nor energy to perform repeated training sessions.

4.3.4 Learning curves for lying posture recognition

Learning curves were plotted to asses the number of training instances required to achieve desired lying posture recognition accuracies. Figure 4.16 depicts the learning curves for each classifier with their best feature subsets up to a cardinality of 5. The figures on the left (4.16a, 4.16c, 4.16e) were plotted using the full class set, while

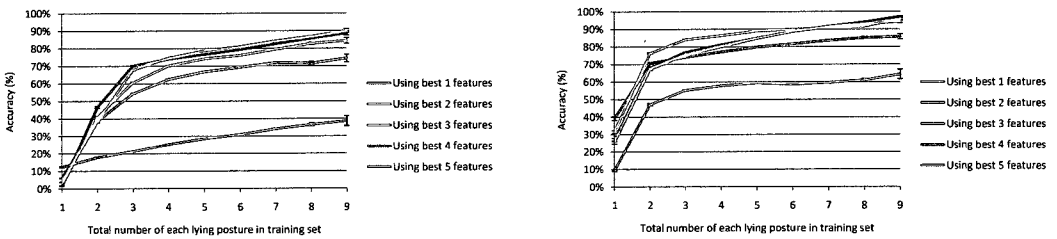

(a) SVM (b) SVM (reduced set of postures)

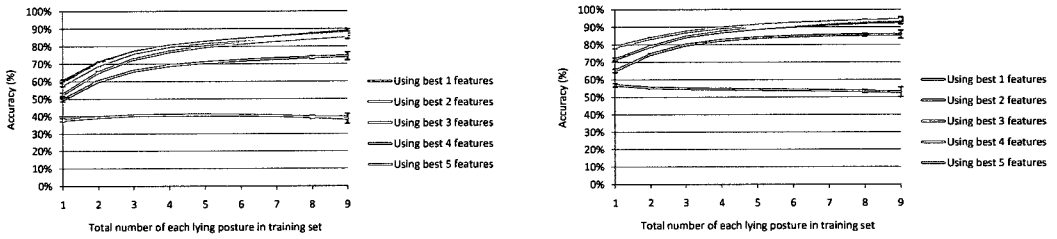

(c) NN (d) NN (reduced set of postures)

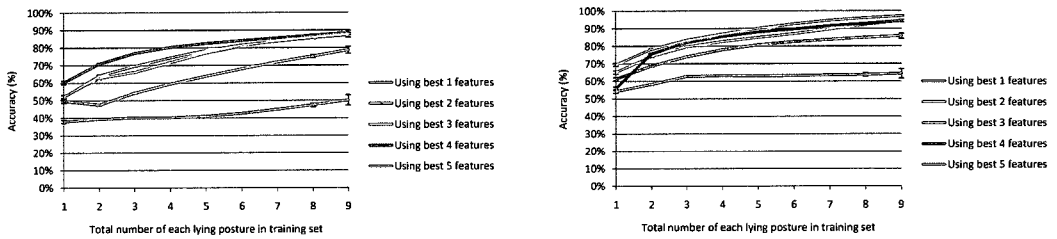

(e) kNN (f) kNN (reduced set of postures)

Figure 4.16: Learning curves for lying posture recognition depicting the total number of each lying posture required in a training set to achieve a particular recognition accuracy.

those on the right (4.16b, 4.16d, 4.16f) were plotted using the reduced set of postures.

With the SVM classifier, performance was very poor with only one example of each posture present in the training set, but increased rapidly with the addition of a second and third training instance for each posture. Performance of the SVM classifier monotonically increased with training set size, however returns in accuracy performance began to diminish once the training set contained at least three examples of each position for the full class set (figure 4.16a) and at least two examples of each position for the reduced class set (figure 4.16b).

The NN classifier's learning rate was more consistent and performance with small training sets was significantly better than with the SVM classifier. Again, performance increased monotonically with increasing training set size, except when only the best single feature was used. With the single best feature, performance remained approximately constant as the training set size was increased.

Not surprisingly, the kNN classifiers learning rates were similar to the NN classifiers. One notable difference in the kNN classifiers performance compared to NN was the appearance of sharp changes in learning rates at particular training set sizes. These sharp changes were the result of a change in the choice of k value by the pre-training parameter optimization procedure.

What may be inferred from these results? It is fairly clear that the SVM classifier is a poor choice when less than three examples of each position are available as training data. Based on the plots, the NN and kNN classifiers would be the appropriate choices in this circumstance. However, to get the best performance from the kNN classifier requires five features and $k = 9$, resulting in a higher computational cost than would be incurred using the NN classifier. The NN classifier only required three features to reach its best performance (and, by definition, $k = 1$). Furthermore, compared to the SVM and kNN classifiers, the NN classifier actually had the best classification performance when only one example of each position was available in the training set.

These results suggest that to begin using the system with a person for which no training data is available, one short training session would be required including a single performance of each lying posture. The NN classifier would be the initial choice as classifier in the lying posture recognizer. Over the course of time, as more training data for that person became available, the SVM classifier could replace NN if the monitoring policy demanded the best possible classification accuracy. Otherwise, the NN classifier would be the overall best choice, considering its excellent performance, low computational cost, and simplicity of implementation.

4.3.5 Comparison to previously published results in lying posture recognition

Table 4.18 shows the lying posture recognition results achieved in the present work and the results achieved in previously published studies. Three key differences distinguish the present work from those listed in the table: the present work 1) uses an optical pressure sensor array instead of an array of force sensitive resistors, 2) uses a larger set of lying postures (seven rather than three or four), and 3) uses a thicker mattress that is typical of a Canadian hospital. With these differences in mind, the results reported in table 4.18 suggest the lying posture recognizer proposed in the present work performs competitively when compared with previously published results. Although the accuracy is below the 100% accuracy reported in [74] and [76], those studies only considered half the number of postures included in the present work, ignoring prone, fetal, or sitting up postures. In addition, it is reasonable to state that the recognition task presented in this work was more challenging because the spatial resolution of the pressure distribution was lower than in the other studies. The inherent spatial resolution of the optical pressure sensor array was limited by the 13 cm inter-sensor distance compared with inter-sensor distances of 5 cm, 7 cm, 7.8 cm, and 10 cm used in the other studies. It is also possible that the thicker mat-

Reference	Sensor configuration	Postures	Mattress configuration	Results
Present work	**132 optical-fibre sensors. 13 cm inter-sensor distance**	**Supine, prone, left/right lateral,left/right fetal,sitting up**	**PressureGuard mattress, 17 cm thick. Sensor array underneath the mattress**	**92.9 ± 2.4% accuracy with full set of postures. 97.1 ± 1.9% accuracy with the reduced class set.**
Hsia et al. [74]	16 long and narrow force sensing resistors (FSRs). 5 cm inter-sensor distance.	Supine, left/right lateral	Unspecified	100% accuracy with all positions parallel to center line. 78.7% average accuracy with different lying angles included.
Seo et al. [75]	336 FSRs. 7 cm vertical spacing between sensors, 5 cm horizontal.	Supine, left/right lateral, sitting up	Sensors embedded at top surface of mattress	93.6% "success rate"
Harada et al. [76]	210 FSRs. 7.8 cm inter-sensor distance	Supine, left/right lateral	Sensors underneath 50 mm thick futon-mat	No quantitative accuracy results reported, but figures suggest 100% accuracy
Nishida et al. [77]	221 FSRs. 10 cm inter-sensor distance at head and legs, otherwise 5 cm inter-sensor distance	Supine, prone, left/right lateral	Sensors underneath 10 cm thick futon-mat	Accuracy measured by number of lying posture transitions, compared with video analysis, reported as a mean difference of 2 transitions.

Table 4.18: Comparison to performance achieved in published research on lying posture recognition using pressure sensor arrays.

tress would contribute to a lower spatial resolution, further spreading the pressure distribution.

Chapter 5

Postural Transition Detection

Figure 5.1 depicts a high-level diagram of the proposed postural transition detection system. Pressure data captured by the pressure sensor array is input to the static activity recognizer and movement detector which provide the necessary movement and posture information needed by the postural transition detector. Note that this thesis focuses exclusively on the lie-to-sit postural transition. The goal of this chapter is to demonstrate a postural transition detector which detects lie-to-sit postural transitions with a low detection miss rate and with an accuracy commensurate with video analysis.

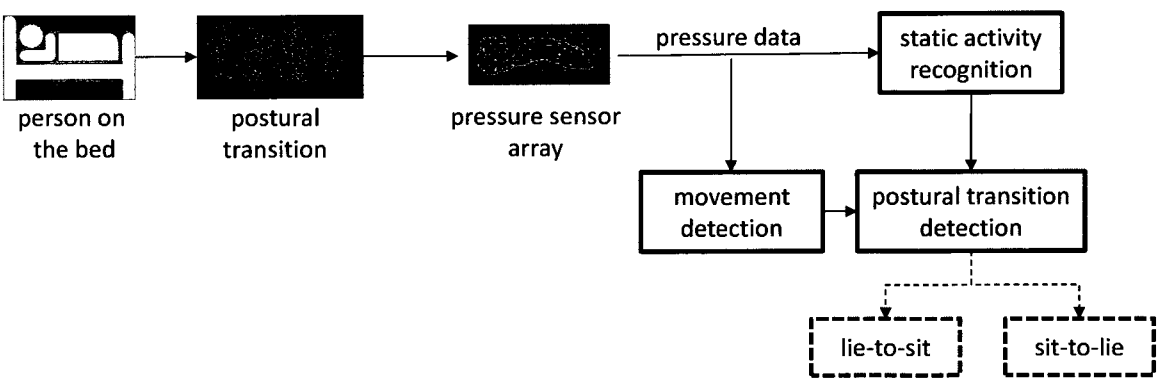

Figure 5.1: Overview of the proposed postural transition detection system.

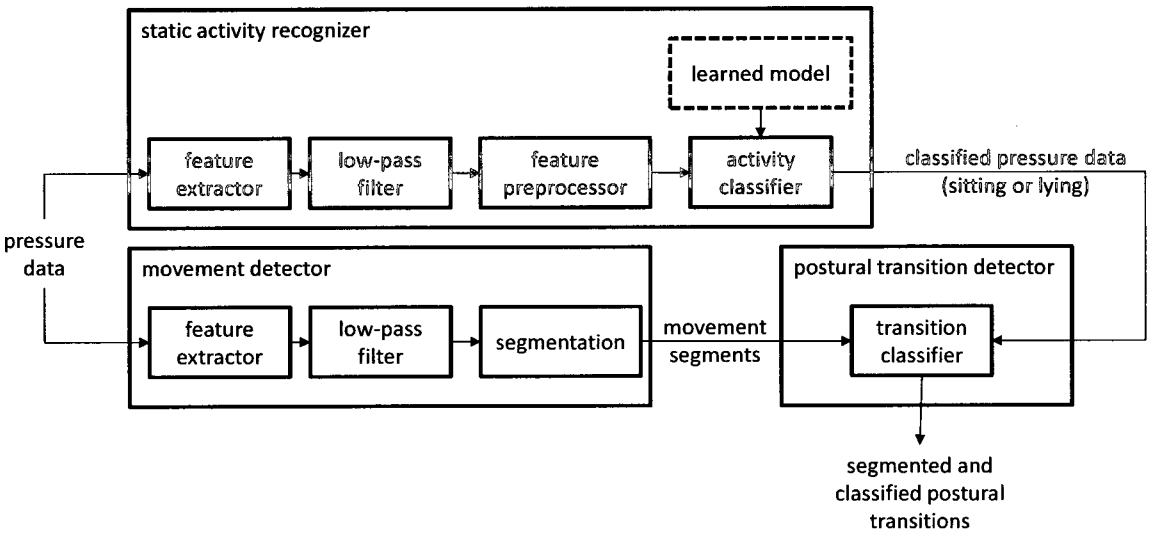

Figure 5.2: Components of the proposed postural transition detection system.

5.1 Methods

Figure 5.2 depicts a block diagram of the components in the postural transition detection system and the subsequent sections explain each component of the system.

5.1.1 Static activity recognizer

The pressure features computed during feature extraction were selected based on the results presented in section 4.3.1. SVM was used as the classifier, as it provided the best estimated static activity recognition performance.

To reduce high-frequency noise in the time sequence of extracted features, an averaging low-pass filter was applied to each feature. In the following discussion, the window length of the averaging low-pass filter is denoted by L_W, and is the same for each feature. The value of L_W was determined using the parameter optimization procedure which is discussed in section 5.2. The rationale for using a low-pass filter was to prevent posture misclassification caused by random movement of the person on the bed.

5.1.2 Movement detector

The static activity recognizer alone would not be sufficient for detecting the duration of a postural transition, as figure 5.3 demonstrates. To overcome the limitation of using a single sample based posture classifier, a movement detection algorithm was implemented.

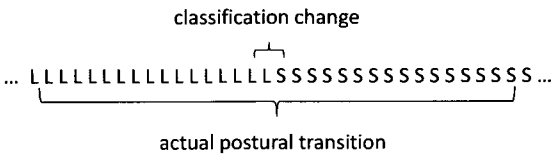

Figure 5.3: Classification change compared with actual postural transition for a sequence of classified pressure samples. L represents a sample classified as 'lying' and S as 'sitting'. This figure depicts a sequence of posture classified samples which contains a postural transition (i.e. lie-to-sit). The change in classification from one sample to the next does not necessarily represent the whole postural transition. As such, more information is required to detect a postural transition.

The movement detector makes one simple assumption: during any subject movement, a measurable change in the center of pressure will occur. Following this assumption, the movement detector simply computes the magnitude of the first backward difference of the center of pressure vector. In other words, this feature is the magnitude of the center of pressure velocity vector. Formally, this feature is expressed as:

$$\Delta M[t] = \sqrt{(\text{LonCP}[t] - \text{LonCP}[t-1])^2 + (\text{LatCP}[t] - \text{LatCP}[t-1])^2} \qquad (5.1)$$

where LonCP$[t]$ and LatCP$[t]$ are as defined in equations 4.7 and 4.8, respectively. $\Delta M[t]$ is hereafter called the *motion vector sequence.*

Values of $\Delta M[t]$ above 0 indicate movement of the center of pressure at time t. The motion vector sequence is segmented using a simple rule based on the observation that with body weight transfers, like those that occur during a postural transition, the value of $\Delta M[t]$ should surpass a threshold value T_M, assuming the transfer occurs

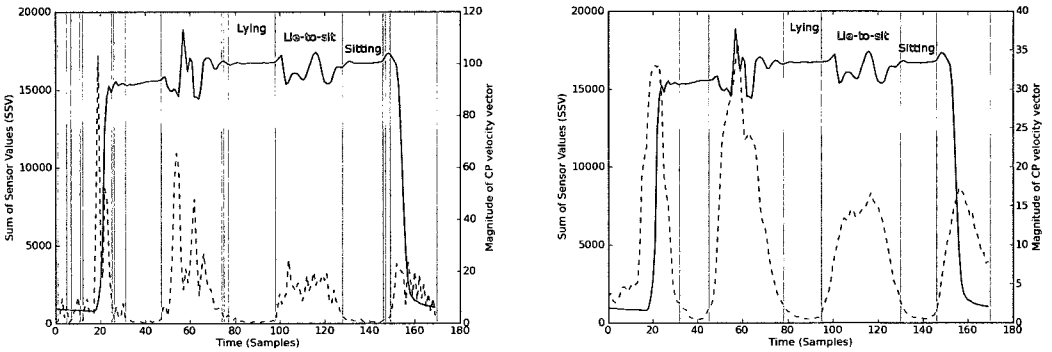

(a) No low-pass filter

(b) With low-pass filter

Figure 5.4: Magnitude of the center of pressure velocity vector for one bed entry and exit routine. The shaded regions show movement segments (i.e. where the magnitude of the center of pressure velocity vector is above the movement threshold T_M)

within a relatively short period of time. Formally, given the motion vector sequence, the movement detector outputs a sequence of movement segments

$$\sigma = \{(t_S, t_E) \mid \Delta M[t_i] > T_M \text{ for all } t_i \in [t_S, t_E]\}. \tag{5.2}$$

The movement detector also includes an optional averaging low-pass filter applied to the motion vector sequence. If the filter is used, the length of the filter window is set equal to the window length used for the feature filter (L_W) in the static activity recognizer. The rationale for the averaging low-pass filter on the motion vector sequence was to further reduce the undesirable effects of random motion. Figure5.4 shows how one bed entry and exit routine is segmented by the movement detector. The dotted line in figure 5.4a shows the magnitude of the center of pressure velocity vector when it is unfiltered and figure 5.4b shows the magnitude of the center of pressure velocity vector when it is low-pass filtered.

5.1.3 Transition classifier

A rule-based transition classifier was used to classify movement segments output by the movement detector. Figure 5.5 depicts the proposed decision tree for the rule-based transition classifier. Let $\Phi[t]$ denote the activity (sitting or lying) recognized in sample t by the static activity recognizer. The following procedure is performed for each segment $(t_S, t_E) \in \sigma$. Let $A = \Phi[t_S]$ and $B = \Phi[t_E]$. First, consider the case where $\Phi[t] \in \{A, B\}$ when $t_S < t < t_E$; in other words, at most two postures are present in the motion sequence. In this case, if $A = B$ then the movement sequence is classified as non-transitional motion, otherwise it is classified as an A-to-B postural transition. The second case to consider is where $\Phi[t] \notin \{A, B\}$ for some $t_S < t < t_E$. In this case, if $A = B$ then the movement is classified as a possibly failed A-to-C transition, otherwise it is classified as an A-to-C transition followed by a C-to-B transition.

5.2 Performance evaluation

Figure 5.6 depicts the performance evaluation approach. A fixed protocol of postural transitions was performed by each person (cf. postural transition data in section 3.4.3). The person's physical activity was recorded for later analysis by both a bed-based optical pressure sensor array and two video cameras. The video data was subject to visual analysis by a medical student and the pressure data for the entire sequence of ten bed entry and exit routines was input to the postural transition detector. Postural transition durations determined by the postural transition detector were then compared to the durations determined via video analysis.

The postural transition detector had two parameters requiring specification: the motion threshold T_M and the low-pass filter window length L_W. Every combination of parameters $T_M \in 0.1, 0.2, \ldots, 4.9, 5.0$ and $L_W \in \{1, 5, 10, 15, 20\}$ was evaluated

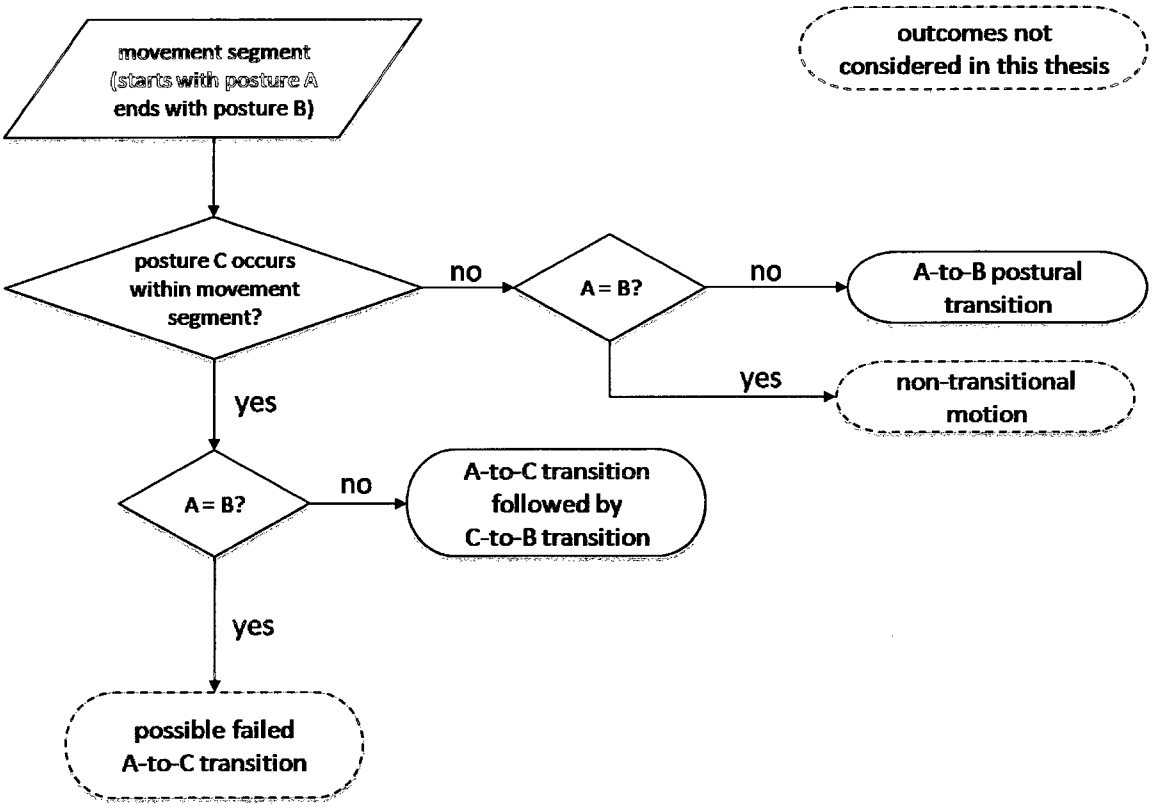

Figure 5.5: Decision tree for the rule-based transition classifier.

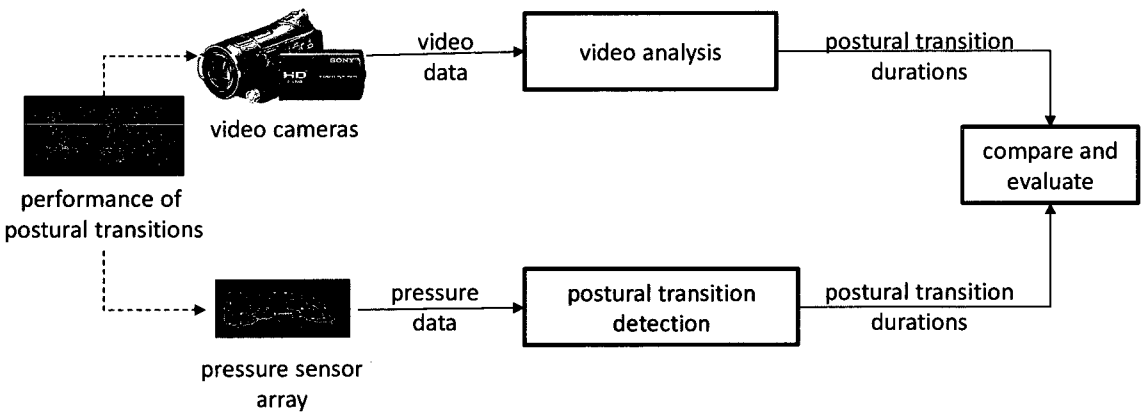

Figure 5.6: Evaluating performance of the postural transition detector.

to determine the optimal parameter values for lie-to-sit detection. The parameters were optimized by minimizing the mean difference between the transition durations reported by the transition detector and those reported by the human video analyst.

To test for significant differences between the postural transition detector and video analysis, two-tailed paired-sample randomization tests (see [101]) were performed using the parameters for which the minimum mean difference between the durations determined with the detector and those determined with video analysis was achieved. A p-value less than 0.05 is typically used to show statistically significant differences, so a p-value greater than 0.05 indicates that a difference is statistically insignificant. Thus, a p-value greater than 0.05 indicates that the accuracy of the postural transition detector is commensurate with video analysis because the differences in recorded durations are statistically insignificant.

5.3 Results and discussion

The performance evaluation was applied to the four participant groups separately followed by a between-group comparison.

5.3.1 Young healthy participant group

Figure 5.7 depicts the results obtained using the data from the young healthy participant group. The upper plot in figure 5.7a shows the mean difference between detector and video analysis durations for all window lengths L_W and thresholds T_M. The plot suggests that the filter window length has little effect on the minimum mean difference with video analysis. The threshold has a larger influence on the mean difference, with an optimal threshold in the range of 0.8-2.0. The minimum mean difference was achieved using $L_W = 15$ and $T_M = 1.4$, with a mean difference of 0.537 ± 0.413 s. The lower plot of Figure 5.7a shows the miss rate of the detector for all window lengths

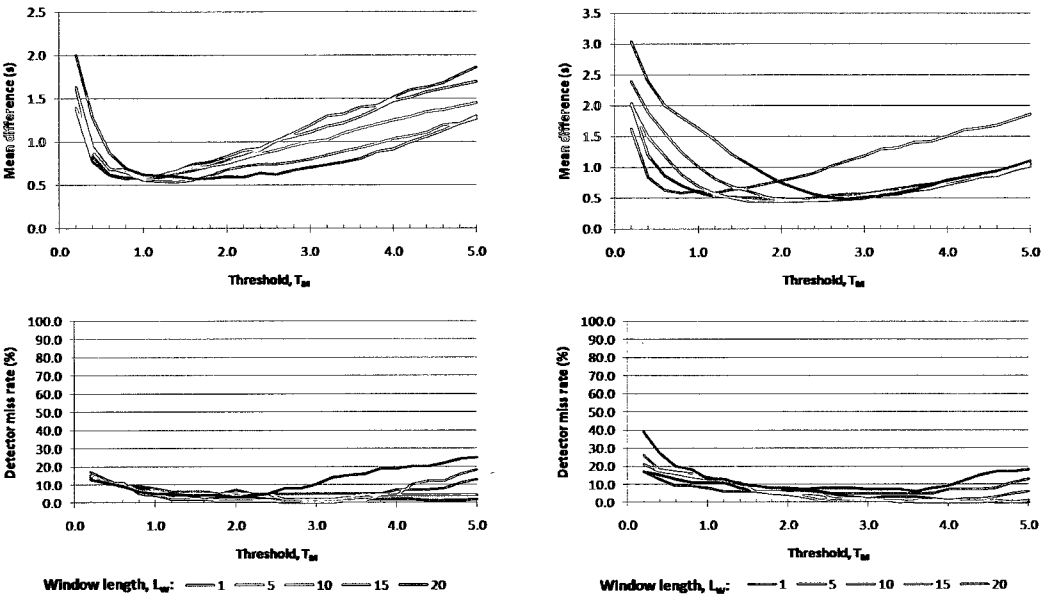

(a) Unfiltered motion vector sequence (b) Averaged motion vector sequence

Figure 5.7: Mean difference between video analysis and postural transition detector reported durations and detector miss rate for all combinations of parameters L_W and T_M with the young healthy participant group.

and thresholds. As with the mean differences, the plot indicates that window length has little effect on the miss rate in the threshold range of interest (0.8-2.0). In this range, the threshold also has only small influence on miss rate, with miss rates falling below 10%. Using the choice of parameters to minimize mean difference ($L_W = 15$ and $T_M = 1.4$), the miss rate was 2%.

The upper plot in Figure 5.7b shows the mean difference between detector and video analysis durations when the motion vector sequence is low-pass filtered (averaged). The choice of filter window length had little effect on the minimum mean difference, however, as the window length increased, the optimal threshold values also increased. The minimum mean difference of 0.430 $\pm$ 0.312 s was achieved using $L_W = 10$ and $T_M = 2.0$. A paired-sample randomization test showed that the differences between video analysis and transition detector when the motion vector was averaged were significantly lower ($p = 0.001$) than the differences resulting from using

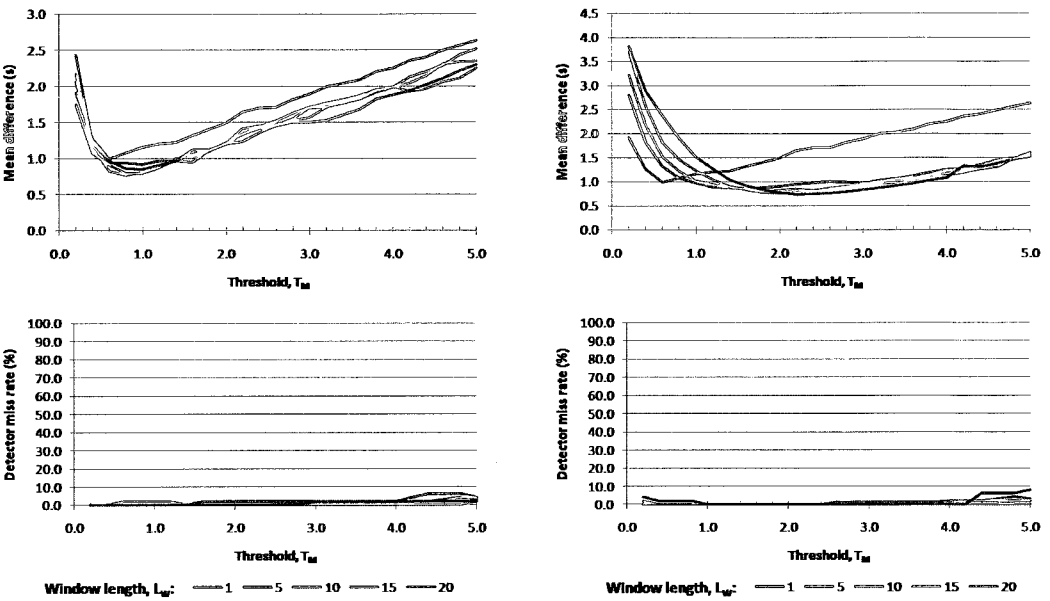

(a) Unfiltered motion vector sequence (b) Averaged motion vector sequence

Figure 5.8: Mean difference between video analysis and postural transition detector reported durations and detector miss rate for all combinations of parameters L_W and T_M with the older healthy participant group.

an unfiltered motion vector. The lower plot of Figure 5.7b shows the miss rate of the detector for all window lengths and thresholds when the motion vector sequence was averaged. The plot demonstrates that the choice of window length has little effect on the miss rate. Using the choice of parameters to minimize mean difference ($L_W = 10$ and $T_M = 2.0$), the miss rate was 4%, which was an increase of 2% over the miss rate obtained without the motion vector filter.

5.3.2 Older healthy participant group

Figure 5.8 shows the results obtained with the older healthy participant group. As with the young healthy participant group, the upper plot of figure 5.9a demonstrates that the filter window length had only a marginal effect on the mean difference between the detector and video analysis durations. With the old healthy participant group, the optimal range for the threshold was approximately 0.6-1.2, with the min-

imum mean difference of 0.763 ± 0.619 s achieved using $L_W = 10$ and $T_M = 0.8$. The lower plot of figure 5.9a depicts the miss rate of the detector on the old healthy participant data. Once again, there was little effect of window length on miss rate. The miss rate was consistently below 10% across all parameters combinations. With the mean difference optimized parameters $L_W = 10$ and $T_M = 0.8$ the miss rate was 2%.

The upper plot of figure 5.9a shows that averaging the motion vector sequence only slightly lowered the minimum mean difference achieved by the detector across parameters. However, a paired-sample randomization test, using $L_W = 20$ and $T_M = 2.2$, showed that the mean differences achieved with an averaged motion vector sequence were not significantly lower than when the motion vector was not averaged ($p = 0.3262$). As with the young healthy participant group, the primary effect of averaging the motion vector sequence was to increase the threshold required to reach minimum mean difference. The minimum mean difference achieved with an averaged motion vector sequence was 0.734 ± 0.515 s with $L_W = 20$ and $T_M = 2.2$. The lower plot of figure 5.9a demonstrates that filtering the motion vector had little effect on the miss rate. This is unsurprising, as the miss rates using an unfiltered motion vector were already very low. With the mean difference optimized parameters $L_W = 20$ and $T_M = 2.2$ the miss rate was 0%, a decrease of 2%.

5.3.3 Hip fracture participant group

Figure 5.9 depicts the results from the hip fracture recovery participant group. Consistent with the young healthy and old healthy groups, the upper plot of figure 5.9a indicates the minimal effect of window length on the resulting mean time difference between the transition detector and video analysis. The lowest mean difference 1.144 ± 0.987 s was achieved using parameters $L_W = 10$ and $T_M = 0.4$. For $T_M > 1.0$, the mean difference between the algorithm and video analysis consistently increased

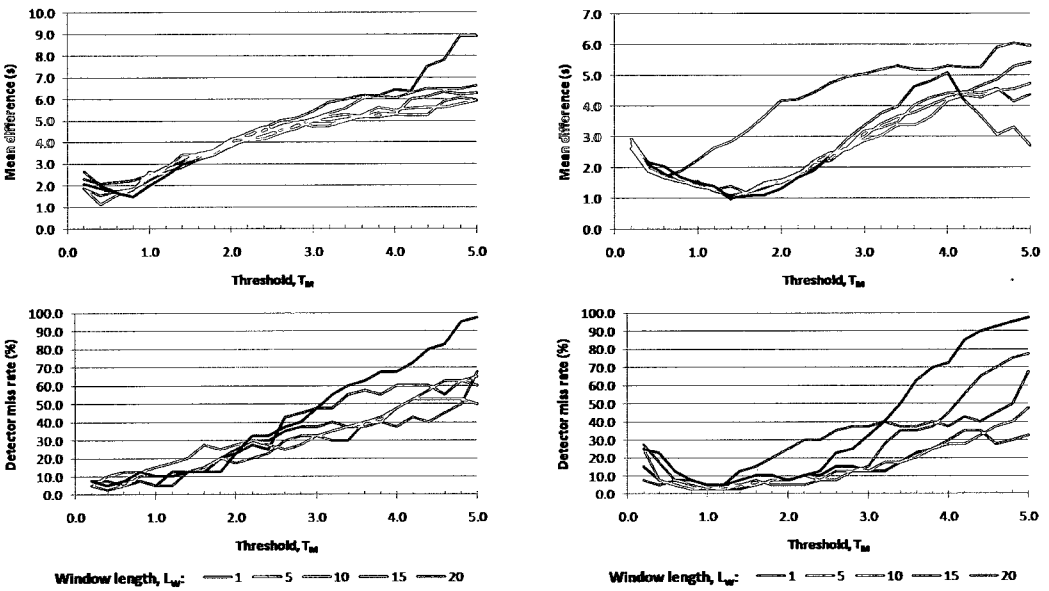

(a) Unfiltered motion vector sequence

(b) Averaged motion vector sequence

Figure 5.9: Mean difference between video analysis and postural transition detector reported durations and detector miss rate for all combinations of parameters L_W and T_M with the hip fracture participant group.

as T_M increased. The lower plot in figure 5.9a shows the miss rate of the transition detector and demonstrates that the lowest miss rates were achieved in the vicinity of the optimal parameter values. As with mean difference, the miss rates began to increase steadily past $T_M = 1.0$ for all window lengths. With $L_W = 10$ and $T_M = 0.4$, 4 out of 40 transitions were missed (10%).

Averaging the motion vector sequence resulted in slight improvements in mean difference, as shown in the top plot of figure 5.9b. As with the young healthy and old healthy participant groups, averaging the motion vector sequence resulted in a shift of the optimal T_M values, with the minimum mean difference of 0.976 ± 1.379 s occurring at parameter values $L_W = 10$ and $T_M = 1.4$. The differences achieved using an averaged motion vector sequence were not significantly less than the differences achieved without averaging ($p = 0.3526$). The bottom plot of figure 5.9b shows the transition detector miss rates when averaging the motion vector sequence. Averaging

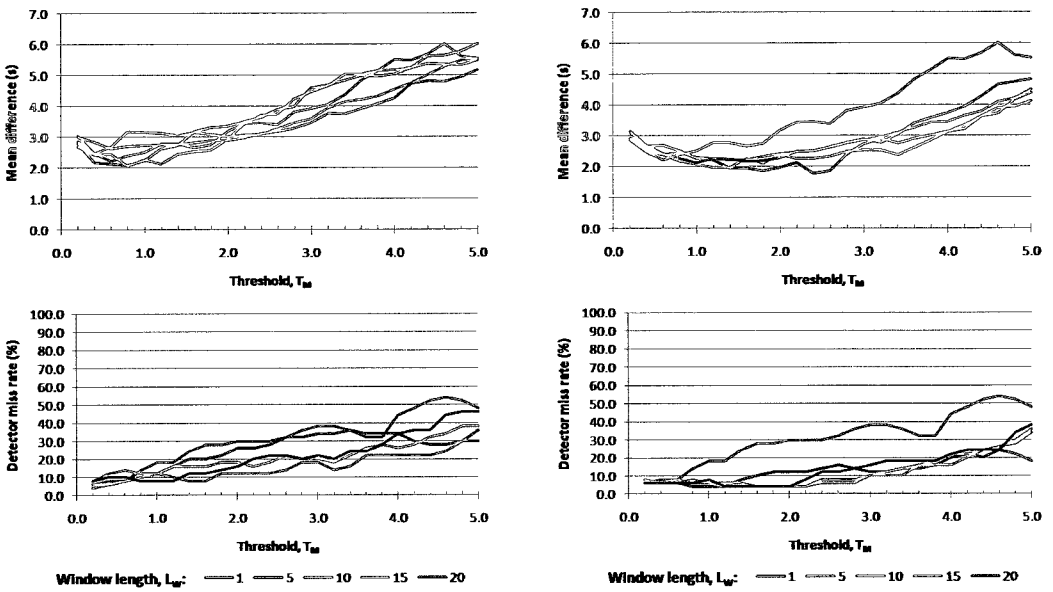

(a) Unfiltered motion vector sequence (b) Averaged motion vector sequence

Figure 5.10: Mean difference between video analysis and postural transition detector reported durations and detector miss rate for all combinations of parameters L_W and T_M with the stroke participant group.

the motion vector sequence resulted in a lower miss rate (5%) at the optimal parameter values $L_W = 10$ and $T_M = 1.4$.

5.3.4 Stroke participant group

Figure 5.10 shows the results obtained with the participants recovering from stroke. The top plot of figure 5.10a shows the mean differences in duration times determined via video analysis and transition detector. As with all the other participant groups, the choice of window length did not have a large effect on the mean difference between video analysis and transition detector. The lowest mean difference between video and transition detector using an unfiltered motion vector was obtained using the parameters $L_W = 10$ and $T_M = 0.8$. At these parameter values, the mean difference was 2.058 ± 1.844 s. The bottom plot of figure 5.10a depicts the miss rates for the transition detector for all tested parameters. Consistent with the other participant

groups, the lowest miss rates were obtained in the vicinity of the optimal parameter values and tended to increase with increasing T_M for all tested values of L_W. At $L_W = 10$ and $T_M = 0.8$ the miss rate was 10% (5 out of 50 transitions were missed).

The top plot of figure 5.10b shows the mean differences between video and transition detector when the motion vector sequence was averaged. Again, averaging the motion vector sequence resulted in slightly lower differences, with the lowest difference of 1.765 ± 1.599 s occurring at parameter values $L_W = 5$ and $T_M = 2.4$. However, the differences achieved with an averaged motion vector were not significantly lower than those achieved without averaging ($p = 0.3604$). The lower plot of figure 5.10b shows the miss rates of the transition detector using an averaged motion vector sequence, and again miss rates were generally lowered. However at the parameter values $L_W = 5$ and $T_M = 2.4$, the miss rate was higher (14%) than the miss rate achieved with the unfiltered motion vector.

5.3.5 Between-group comparison

The results obtained from the postural transition detector with each participant group suggests that postural transition detection was more easily accomplished with the young healthy and older healthy participant groups. The top plot of figure 5.11a depicts the mean differences in measured transition duration between the transition detector and video analysis using the value of L_W which resulted in the lowest mean difference and an unfiltered motion vector sequence across all tested values of T_M. The plot indicates that the closest match between the transition detector and video analysis, as measured by mean difference, occurred with the young healthy participant group. Furthermore, the standard errors of the mean were, in general, lowest with the young healthy group, suggesting that the transition detector was most consistently agreeing with video analysis for the young healthy group. After the young healthy group, the transition detector most agreed with video analysis on the older healthy

group. Compared with the hip fracture and stroke participant groups, the transition detector had a lower mean difference with video analysis and smaller standard errors when detecting transitions with the older healthy participant group. Standard errors were greater in the hip fracture and stroke participant groups, suggesting that differences between the transition detector and video analysis were greater with these groups.

The bottom plot of figure 5.11a shows the miss rates for the transition detector for each participant group. Overall, miss rates were lowest with the healthy participant groups than with the stroke and hip fracture groups, suggesting that transitions were more easily captured in the healthy groups. The miss rate for the young healthy group with optimal parameter values was 2%, for the older healthy group 2%, for the hip fracture group 10%, and for the stroke group 10%.

Figure 5.11b shows the mean differences and miss rates for the transition detector across groups using the averaged motion vector sequence. The overall effect of low-pass filtering the motion vector sequence was to lower the mean differences between transition detector and video analysis. As the top plot of figure 5.11b shows, the effect was most pronounced with the hip fracture and stroke groups. With the young healthy and older healthy groups, the effect of an averaged motion vector sequence was less pronounced.

The bottom plot of figure 5.11b demonstrates graphically that the miss rates for the stroke and hip fracture groups were significantly reduced, while with the healthy groups the miss rates were not significantly affected. A paired-sample randomization test confirmed this observation statistically, showing significant reductions in miss rate for both the hip fracture ($p = 0.0480$) and stroke ($p = 0.0006$) groups, but no significant reductions for the healthy groups ($p = 0.4948$ for young healthy and $p = 0.9800$ for older healthy).

Table 5.1 summarizes the optimized parameter values for each group, the detec-

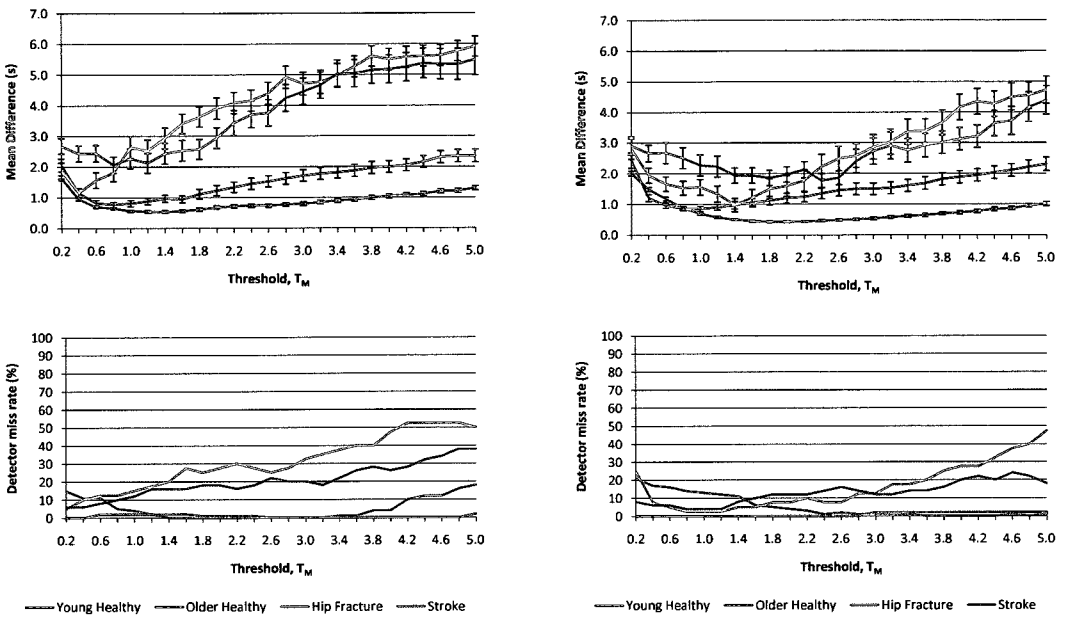

(a) Unfiltered motion vector

(b) Averaged motion vector

Figure 5.11: Comparison of sit-to-lie postural transition detection results between all participant groups, using the best parameters for each group.

tion miss rates, and reports the result of a paired-sample randomization test which tests whether there was a significant difference between the postural transition detector and video analysis. Detection miss rate refers to when a lie-to-sit postural transition was not detected at all in a bed entry and exit routine. Misses were typically caused by short sitting phases, so the movement detector did not detect a state of rest. Table 5.1a demonstrates that for all groups except the young healthy (with a 0.05 significance level), the differences between the transition detector and video analysis when the motion vector sequence was not averaged were statistically insignificant. Table 5.1b reports that with the addition of an averaging filter on the motion vector sequence, the transition detector reported lie-to-sit transition durations with statistically insignificant differences from the durations reported by video analysis for all participant groups.

The detection miss rates were within reasonable limits. If the postural transition detector was used to monitor long-term changes in mobility, which is one of the main

Participant group	L_W	T_M	Detection miss rate (%)	Paired-sample randomization test of significant differences between video and detector (p-value)
Young healthy	15	1.4	2	0.0204
Older healthy	10	0.8	2	0.7392
Hip fracture	10	0.4	10	0.1552
Stroke	10	0.8	10	0.8448

(a) Unfiltered motion vector

Participant group	L_W	T_M	Detection miss rate (%)	Paired-sample randomization test of significant differences between video and detector (p-value)
Young Healthy	10	2.0	4	0.8894
Older healthy	20	2.2	0	0.9024
Hip fracture	10	1.4	5	0.5608
Stroke	5	2.4	14	0.0960

(b) Averaged motion vector

Table 5.1: Summary of best parameters, best detection miss rate, and results of paired-sample randomization tests of the significance of differences between durations reported by video analysis and postural transition detector. Detection miss rate refers to when a lie-to-sit postural transition was not detected at all in a bed entry and exit routine. Misses were typically caused by short sitting phases, so the movement detector did not detect a state of rest.

motivations for such a system, missing one in ten (10%) lie-to-sit postural transitions would not adversely affect the systems ability to monitor trends in a persons mobility because such trends would vary slowly when compared to the frequency of lie-to-sit postural transitions in daily life.

The aim of this chapter was to demonstrate that an optical pressure sensor array, placed beneath a typical hospital mattress, can detect postural transitions with low miss rate and report the duration of postural transitions with an accuracy commensurate with a medical student determining durations via video analysis. In general, the

results suggest that the postural transition detector could be used to detect the lie-to-sit transitions of young healthy, older healthy, hip fracture, and stroke patients, and that the durations reported by the detector would differ insignificantly from durations reported using video analysis.

Chapter 6

Conclusions

This work investigated the two related subjects of posture recognition and postural transition detection using optical bed-based pressure sensor arrays, focusing specifically on recognizing lying and sitting activities, lying postures, and detecting lie-to-sit postural transitions. Using a mattress typical of a Canadian hospital, one goal was to implement static activity and lying posture recognizers for an optical bed-based pressure sensor array which could compete with the performance reported in previously published results using both wearable sensors or other pressure sensitive mat technology. The other goal was to implement a lie-to-sit postural transition detector with a low detection miss rate and an accuracy commensurate with video analysis. This chapter begins with a review of how the objectives expressed in the introductory section 1.3 were achieved and the contributions to knowledge provided by this work. The chapter ends with a discussion of potential avenues for future research.

6.1 Review of objectives and research question

The ten methodological objectives set out in section 1.3 to address the research question were achieved, and can be summarized as follows:

1. eight candidate pressure signal features for posture recognition were identified in section 4.1.1: the sum of sensor values (SSV), the number of active sensors (NAS), the weighted sum of sensor values (WSSV), the weighted number of active sensors (WNAS), the longitudinal center of pressure (LonCP), the lateral center of pressure (LatCP), the longitudinal variance (LonV), and the lateral variance (LatV);

2. three candidate classification techniques were identified in section 4.1.2: support vector machine (SVM), nearest neighbour (NN), and k-nearest neighbour (kNN);

3. section 4.3.1 demonstrated the following:

   (a) the combination of WNAS, LatCP, and LonV candidate features with the SVM classifier resulted in the best subject-dependent static activity recognition performance;

   (b) the combination of the NAS, WSSV, LatCP, LatV, and LonV candidate features with the SVM classifier resulted in the best subject-independent static activity recognition performance;

4. section 4.3.2 compared performance of the static activity recognizer with previously published results and demonstrated that the bed-based optical pressure sensor array achieved performance competitive with wearable sensor technologies for recognizing sitting and lying activities when a person is on the bed;

5. section 4.3.3 demonstrated the following:

   (a) the combination of all eight candidate features with the SVM classifier resulted in the best lying posture recognition performance, however the best performance when the accuracy to computation ratio was considered was achieved using the SSV, WNAS, LonCP, and LonV features with SVM;

(b) the combination of SSV, WSSV, LatCP, and LonCP with the SVM classifier resulted in the best lying posture recognition performance when the left/right lateral and left/right fetal positions were combined into lateral and fetal classes, respectively;

6. section 4.3.5 demonstrated that the lying posture recognizers performance was competitive with previously published results using force sensitive resistor based pressure sensor arrays;

7. section 5.1.2 described the implementation of a simple movement detection algorithm;

8. a rule-based postural transition classifier, combining inputs from a static activity recognizer and movement detector, was presented in section 5.1.3;

9. section 5.3.5 showed that an effective set of system parameters were found to reliably detect lie-to-sit postural transitions in the young healthy, older healthy, stroke, and hip fracture participant groups; and

10. section 5.3.5 also showed that the postural transition detection system could detect lie-to-sit postural transitions with reasonably low miss rates and report the duration of lie-to-sit postural transitions with statistically insignificant differences when compared with the durations reported by a medical student using video analysis.

All three aspects of the research question that were set out in section 1.1 were addressed. Using a bed-based pressure sensor array and the static activity recognizer described in chapter 4, a persons sitting and lying static activities while they are on a bed can be automatically recognized with the same reliability afforded by employing wearable sensors. A complete set of lying postures may also be reliably and automatically recognized using a bed-based pressure sensor array placed underneath

a mattress typical of Canadian hospitals. Finally, lie-to-sit postural transitions may be detected with acceptable reliability and their durations reported with an accuracy commensurate of a medical student analyzing video recordings of the same postural transitions.

The unobtrusiveness of the bed-based pressure sensor array follows from the fact that it is placed beneath a mattress and monitors bed-based activities, so there is no hardware that must be attached to the body. The reliability of the static activity and lying posture recognizers was established by comparison to previously published results using both wearable sensors and different pressure sensitive array technology.

6.2 Review of contributions

Early results with the postural transition detector were used to

1. demonstrate that a bed-based pressure sensor array can detect differences in postural transition duration between young healthy, older healthy, and older stroke patients;

    - [28] P. Carlson, N. Foubert, F. Knoefel, R. Goubran, M. Bilodeau, and H. Sveistrup, "Smart mat technology: Can it differentiate bed transfers in young healthy, old healthy and old stroke populations?" in *Canadian Geriatrics Society Annual General Meeting*, Toronto, Canada, April 2009

2. demonstrate that a bed-based pressure sensor array can detect differences in postural transition duration between older healthy and older hip fracture patients.

    - [29] S. Mondoux, P. Carlson, N. Foubert, A. Arcelus, F. Knoefel, R. Goubran, M. Bilodeau, and H. Sveistrup, "Smart mat technology: Can

it differentiate bed transfers in older healthy and older hip fracture patients?" in *Regional Geriatrics Program Annual General Meeting*, Ottawa, Canada, October 2009.

In addition, this thesis contributed new knowledge by:

3. demonstrating, using simple pressure signal features, that bed-based pressure sensor arrays can be used to reliably recognize lateral decubitus, prone, supine, and fetal lying postures when placed beneath a mattress typical of Canadian hospitals;

4. providing the first analysis of lie-to-sit postural transition detection using bed-based pressure sensor arrays;

5. demonstrating that bed-based pressure sensor arrays can be used to unobtrusively detect and determine the durations of lie-to-sit postural transitions with an accuracy commensurate of a medical student analyzing video recordings of the same postural transitions.

6. determining a set of simple pressure signal features such that bed-based pressure sensor arrays may be used to reliably recognize sitting and lying static activities when performed in a variety of locations and orientations on a bed;

7. demonstrating that a nearest neighbour classifier is a better choice of pattern classifier than a support vector machine for recognizing lying postures when limited training data is available; and

8. discovering that an effective fetal lying posture recognizer could be implemented using a simple nearest neighbour classification of a persons longitudinal center of pressure on the mattress.

6.3 Future research

To conclude, we suggest several ways in which this work could be extended in future research. The most difficult task for the lying posture recognizer discussed in chapter 4 was discriminating between supine and prone positions. Future work could investigate different pressure signal features with the aim of improving the ability to distinguish supine and prone lying postures. The lying posture recognizer implemented in this thesis only considered static pressure features (i.e. without considering how they change over time). An investigation into the dynamics of the proposed features during periods in which the person transitions between prone and supine postures may perhaps reveal a method to better distinguish prone and supine.

Although the lie-to-sit postural transition detector was able to detect transitions with reasonable miss rates, these rates could likely be lowered with more a more sophisticated movement detection scheme. In the same vein, implementation of the rule-based motion classifier depicted in figure 5.5 could be completed such that the system recognizes failed postural transitions and non-transitional movements.

Chapter 5 focused exclusively on lie-to-sit postural transitions, but the methods presented could be extended to recognize sit-to-lie postural transitions and also to recognize transitions between different lying postures.

Relatively simple descriptive statistics were used as features. These features were shown to be sufficient for reliably detecting sitting and lying activities and lying postures, but the extent to which these features could be used for more sophisticated bed-based posture recognition tasks is unknown. Additional work should be carried out to determine whether these features could, for instance, accurately determine the location and orientation of a person on the bed. One way in this information would be useful is for helping prevent people from falling out of bed. Furthermore, this thesis did not investigate the recognition of particular sitting postures; it would be interesting to see if the features evaluated could also be used to recognize different

types of sitting.

The pressure sensor array we used had a lower spatial resolution than most of the other sensor arrays employed in lying posture recognition research. It would be fruitful to determine the smallest number of sensors which could be used to reliably determine sitting and lying postures. In transitioning from a research system to one which could be used in hospitals and long-term care facilities the cost of such systems would become increasingly important. Using sensor arrays with fewer sensors would result in lower manufacturing costs and in turn lower retail costs for such systems.

Finally, the results of this work could be integrated into a complete system for remotely and unobtrusively reporting the physical context of people in bed, including factors such as assumed posture, postural transitions, breathing, restlessness, bed entry and exit, and others.

References

[1] S. Carstairs and W. Keon, "Canada's aging population: Seizing the opportunity. special senate committee on aging: Final report," 2009.

[2] C. Scanaill, B. Ahearne, and G. Lyons, "Long-term telemonitoring of mobility trends of elderly people using SMS messaging," *IEEE Transactions on Information Technology in Biomedicine*, vol. 10, no. 2, p. 412, 2006.

[3] M. Jones, A. Arcelus, R. Goubran, and F. Knoefel, "A pressure sensitive home environment," in *IEEE International Workshop on Haptic Audio Visual Environments and their Applications (HAVE 2006)*, 2006, pp. 10–14.

[4] M. Jones, R. Goubran, and F. Knoefel, "Reliable respiratory rate estimation from a bed pressure array," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 1, 2006, p. 6410.

[5] M. Mathie, A. Coster, N. Lovell, and B. Celler, "Accelerometry: providing an integrated, practical method for long-term, ambulatory monitoring of human movement," *Physiological Measurement*, vol. 25, no. 2, pp. 1–20, 2004.

[6] U. Varshney, "Pervasive healthcare and wireless health monitoring," *Mobile Networks and Applications*, vol. 12, no. 2, pp. 113–127, 2007.

[7] W. Wong, M. Wong, and K. Lo, "Clinical applications of sensors for human

posture and movement analysis: A review," *Prosthetics and Orthotics International*, vol. 31, no. 1, pp. 62–75, 2007.

[8] A. Neill, S. Angus, D. Sajkov, and R. McEvoy, "Effects of sleep posture on upper airway stability in patients with obstructive sleep apnea," *American Journal of Respiratory and Critical Care Medicine*, vol. 155, no. 1, p. 199, 1997.

[9] R. Cartwright, F. Diaz, and S. Lloyd, "The effects of sleep posture and sleep stage on apnea frequency," *Sleep*, vol. 14, no. 4, p. 351, 1991.

[10] R. McEvoy, D. Sharp, and A. Thornton, "The effects of posture on obstructive sleep apnea," *The American Review of Respiratory Disease*, vol. 133, no. 4, p. 662, 1986.

[11] Y. Itasaka, S. Miyazaki, K. Ishikawa, and K. Togawa, "The influence of sleep position and obesity on sleep apnea," *Psychiatry and Clinical Neurosciences*, vol. 54, no. 3, pp. 340–341, 2000.

[12] S. Martin, I. Marshall, and N. Douglas, "The effect of posture on airway caliber with the sleep-apnea/hypopnea syndrome," *American Journal of Respiratory and Critical Care Medicine*, vol. 152, no. 2, p. 721, 1995.

[13] Y. Matsuzawa, S. Hayashi, S. Yamaguchi, S. Yoshikawa, K. Okada, K. Fujimoto, and M. Sekiguchi, "Effect of prone position on apnea severity in obstructive sleep apnea," *Internal Medicine*, vol. 34, no. 12, pp. 1190–1193, 1995.

[14] N. Yildirim, M. Fitzpatrick, K. Whyte, R. Jalleh, A. Wightman, and N. Douglas, "The effect of posture on upper airway dimensions in normal subjects and in patients with the sleep apnea/hypopnea syndrome," *The American Review of Respiratory Disease*, vol. 144, no. 4, p. 845, 1991.

[15] I. Szollosi, T. Roebuck, B. Thompson, and M. Naughton, "Lateral sleeping position reduces severity of central sleep apnea/Cheyne-Stokes respiration," *Sleep*, vol. 29, no. 8, p. 1045, 2006.

[16] H. Kinney and B. Thach, "The sudden infant death syndrome," *New England Journal of Medicine*, vol. 361, no. 8, p. 795, 2009.

[17] M. Pasquale-Styles, P. Tackitt, and C. Schmidt, "Infant death scene investigation and the assessment of potential risk factors for asphyxia: a review of 209 sudden unexpected infant deaths," *Journal of Forensic Sciences*, vol. 52, no. 4, p. 924, 2007.

[18] M. Willinger, H. Hoffman, and R. Hartford, "Infant sleep position and risk for sudden infant death syndrome," *Pediatrics*, vol. 93, no. 5, pp. 814–819, 1994.

[19] J. Tobin, P. McCloud, and D. Cameron, "Posture and gastro-oesophageal reflux: a case for left lateral positioning," *British Medical Journal*, vol. 76, no. 3, p. 254, 1997.

[20] R. Khoury, L. Camacho-Lobato, P. Katz, M. Mohiuddin, and D. Castell, "Influence of spontaneous sleep positions on nighttime recumbent reflux in patients with gastroesophageal reflux disease," *The American Journal of Gastroenterology*, vol. 94, no. 8, pp. 2069–2073, 1999.

[21] D. Hofsøy, J. Clauss, and B. Wolf, "An intelligent implant system for monitoring and biofeedback therapy of snoring," in *World Congress on Medical Physics and Biomedical Engineering.* Munich, Germany: Springer, September 2009, pp. 196–199.

[22] J. Carr, R. Shepherd, L. Nordholm, and D. Lynne, "Investigation of a new motor assessment scale for stroke patients," *Physical Therapy*, vol. 65, no. 2, p. 175, 1985.

[23] J. Simondson, P. Goldie, and K. Greenwood, "The mobility scale for acute stroke patients: concurrent validity," *Clinical rehabilitation*, vol. 17, no. 5, p. 558, 2003.

[24] J. Simondson, P. Goldie, K. Brock, and J. Nosworthy, "The mobility scale for acute stroke patients: intra-rater and inter-rater reliability," *Clinical Rehabilitation*, vol. 10, no. 4, p. 295, 1996.

[25] M. Rahimpour, N. Lovell, B. Celler, and J. McCormick, "Patients' perceptions of a home telecare system," *International Journal of Medical Informatics*, vol. 77, no. 7, pp. 486–498, 2008.

[26] G. Demiris, M. Rantz, M. Aud, K. Marek, H. Tyrer, M. Skubic, and A. Hussam, "Older adults' attitudes towards and perceptions of'smart home'technologies: a pilot study," *Informatics for Health and Social Care*, vol. 29, no. 2, pp. 87–94, 2004.

[27] G. Demiris, B. Hensel, M. Skubic, and M. Rantz, "Senior residents' perceived need of and preferences for "smart home" sensor technologies," *International Journal of Technology Assessment in Health Care*, vol. 24, no. 01, pp. 120–124, 2008.

[28] P. Carlson, N. Foubert, F. Knoefel, R. Goubran, M. Bilodeau, and H. Sveistrup, "Smart mat technology: Can it differentiate bed transfers in young healthy, old healthy and old stroke populations?" in *Canadian Geriatrics Society Annual General Meeting*, Toronto, Canada, April 2009.

[29] S. Mondoux, P. Carlson, N. Foubert, A. Arcelus, F. Knoefel, R. Goubran, M. Bilodeau, and H. Sveistrup, "Smart mat technology: Can it differentiate bed transfers in older healthy and older hip fracture patients?" in *Regional Geriatrics Program Annual General Meeting*, Ottawa, Canada, October 2009.

[30] R. Duda, P. Hart, and D. Stork, *Pattern classification.* Wiley-Interscience, 2001.

[31] C. Burges, "A tutorial on support vector machines for pattern recognition," *Data Mining and Knowledge Discovery*, vol. 2, no. 2, pp. 121–167, 1998.

[32] J. Platt, "Using analytic QP and sparseness to speed training of support vector machines," *Advances in Neural Information Processing Systems*, pp. 557–563, 1999.

[33] C.-W. Hsu, C.-C. Chang, and C.-J. Lin, "A practical guide to support vector classification," Department of Computer Science, National Taiwan University, 2009. [Online]. Available: http://www.csie.ntu.edu.tw/ cjlin/papers/guide/guide.pdf

[34] R. Begg, M. Palaniswami, and B. Owen, "Support vector machines for automated gait classification," *IEEE Transactions on Biomedical Engineering*, vol. 52, no. 5, pp. 828–838, 2005.

[35] M. Oskoei and H. Hu, "Support vector machine-based classification scheme for myoelectric control applied to upper limb," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 8, pp. 1956–1965, 2008.

[36] N. Saenz-Lechon, V. Osama-Ruiz, J. Godino-Llorente, M. Blanco-Velasco, F. Cruz-Roldan, and J. Arias-Londono, "Effects of audio compression in automatic detection of voice pathologies," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 12, p. 2831, 2008.

[37] H. Narasimha-Iyer, J. Beach, B. Khoobehi, and B. Roysam, "Automatic identification of retinal arteries and veins from dual-wavelength images using structural and functional features," *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 8, pp. 1427–1435, 2007.

[38] S. Osowski, L. Hoai, and T. Markiewicz, "Support vector machine-based expert system for reliable heartbeat recognition," *IEEE Transactions on Biomedical Engineering*, vol. 51, no. 4, pp. 582–589, 2004.

[39] N. Ravi, N. Dandekar, P. Mysore, and M. Littman, "Activity recognition from accelerometer data," in *National Conference on Artificial Intelligence*, vol. 20, no. 3, 2005, p. 1541.

[40] T. Cover and P. Hart, "Nearest neighbor pattern classification," *IEEE Transactions on Information Theory*, vol. 13, no. 1, pp. 21–27, 1967.

[41] "k nearest neighbours," Orange Website, available at http://www.ailab.si/orange/doc/reference/kNNLearner.htm.

[42] S. Bay, "Nearest neighbor classification from multiple feature subsets," *Intelligent Data Analysis*, vol. 3, no. 3, pp. 191–209, 1999.

[43] N. Mezghani, S. Husse, K. Boivin, K. Turcot, R. Aissaoui, N. Hagemeister, and J. de Guise, "Automatic classification of asymptomatic and osteoarthritis knee gait patterns using kinematic data features and the nearest neighbor classifier," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 3, pp. 1230–1232, 2008.

[44] H. Firpi, E. Goodman, and J. Echauz, "Epileptic seizure detection using genetically programmed artificial features," *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 2, pp. 212–224, 2007.

[45] R. Jafari, W. Li, R. Bajcsy, S. Glaser, and S. Sastry, "Physical activity monitoring for assisted living at home," in *4th International Workshop on Wearable and Implantable Body Sensor Networks (Bsn 2007)*, vol. 13. Springer, March 2007, p. 213.

[46] "Tafeta: Smart systems for health," TAFETA Website, http://tafeta.ca.

[47] M. Holtzman, A. Arcelus, R. Goubran, and F. Knoefel, "Breathing signal fusion in pressure sensor arrays," in *IEEE International Workshop on Medical Measurements and Applications (MeMeA 2008)*, 2008, pp. 71–76.

[48] M. Jones, R. Goubran, and F. Knoefel, "Identifying movement onset times for a bed-based pressure sensor array," in *IEEE International Workshop on Medical Measurement and Applications (MeMea 2006)*. IEEE Computer Society Washington, DC, USA, 2006, pp. 111–114.

[49] D. Townsend, R. Goubran, M. Frize, and F. Knoefel, "Preliminary results on the effect of sensor position on unobtrusive rollover detection for sleep monitoring in smart homes," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 1, 2009, p. 6135.

[50] D. Townsend, M. Holtzman, R. Goubran, M. Frize, and F. Knoefel, "Simulated central apnea detection using the pressure variance," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 1, 2009, p. 3917.

[51] A. Arcelus, M. Holtzman, R. Goubran, H. Sveistrup, P. Guitard, and F. Knoefel, "Analysis of commode grab bar usage for the monitoring of older adults in the smart home environment." in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 1, 2009, p. 6155.

[52] H. Tan, L. Slivovsky, and A. Pentland, "A sensing chair using pressure distribution sensors," *IEEE/ASME Transactions On Mechatronics*, vol. 6, no. 3, pp. 261–268, 2001.

[53] L. Slivovsky and H. Tan, "A real-time sitting posture tracking system," in *9th*

International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, vol. 69, 2000, pp. 1049–1056.

[54] H. Morishita, R. Fukui, and T. Sato, "High resolution pressure sensor distributed floor for future human-robot symbiosis environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2002)*, Lausanne, Switzerland, 2002, pp. 1246–1251.

[55] M. Addlesee, A. Jones, F. Livesey, and F. Samaria, "The ORL active floor," *IEEE Personal Communications*, vol. 4, pp. 35–41, 1997.

[56] S. Pirttikangas, J. Suutala, J. Riekki, and J. Röning, "Footstep identification from pressure signals using Hidden Markov Models," in *Finnish Signal Processing Symposium (FINSIG03)*. Citeseer, 2003, pp. 124–128.

[57] ——, "Learning vector quantization in footstep identification," in *3rd IASTED International Conference on Artificial Intelligence and Applications (AIA 2003)*, 2003, pp. 413–417.

[58] K. Koho, J. Suutala, T. Seppänen, and J. Röning, "Footstep pattern matching from pressure signals using segmental semi-markov models," in *12th European Signal Processing Conference (EUSIPCO 2004)*. Citeseer, September 2004, pp. 6–10.

[59] J. Suutala and J. Röning, "Methods for person identification on a pressure-sensitive floor: Experiments with multiple classifiers and reject option," *Information Fusion*, vol. 9, no. 1, pp. 21–40, 2008.

[60] L. Middleton, A. Buss, A. Bazin, and M. Nixon, "A floor sensor system for gait recognition," in *Proceedings of the Fourth IEEE Workshop on Automatic Identification Advanced Technologies*. Citeseer, 2005, pp. 171–176.

[61] A. Dey, G. Abowd, and D. Salber, "A context-based infrastructure for smart environments," in *1st International Workshop on Managing Interactions in Smart Environments (MANSE99)*. Citeseer, 1999, pp. 114–128.

[62] B. Schilit, N. Adams, R. Want *et al.*, "Context-aware computing applications," in *Workshop on Mobile Computing Systems and Applications*. Citeseer, 1994, pp. 85–90.

[63] J. Baek, G. Lee, W. Park, and B. Yun, "Accelerometer signal processing for user activity detection," *Lecture Notes in Computer Science*, pp. 610–617, 2004.

[64] S. Lee, H. Park, S. Hong, K. Lee, and Y. Kim, "A study on the activity classification using a triaxial accelerometer," in *25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 3, 2003.

[65] M. Ermes, J. Parkka, J. Mantyjarvi, and I. Korhonen, "Detection of daily activities and sports with wearable sensors in controlled and uncontrolled conditions," *IEEE Transactions on Information Technology in Biomedicine*, vol. 12, no. 1, pp. 20–26, 2008.

[66] K. Culhane, G. Lyons, D. Hilton, P. Grace, and D. Lyons, "Long-term mobility monitoring of older adults using accelerometers in a clinical environment," *Clinical Rehabilitation*, vol. 18, no. 3, p. 335, 2004.

[67] G. Lyons, K. Culhane, D. Hilton, P. Grace, and D. Lyons, "A description of an accelerometer-based mobility monitoring technique," *Medical Engineering and Physics*, vol. 27, no. 6, pp. 497–504, 2005.

[68] E. Tapia, S. Intille, W. Haskell, K. Larson, J. Wright, A. King, and R. Friedman, "Real-time recognition of physical activities and their intensities using wireless accelerometers and a heart rate monitor," in *International Symposium on Wearable Computers*. Citeseer, 2007.

[69] B. Najafi, K. Aminian, F. Loew, Y. Blanc, and P. Robert, "Measurement of stand-sit and sit-stand transitions using a miniature gyroscope and its application in fall risk evaluation in the elderly," *IEEE Transactions on Biomedical Engineering*, vol. 49, no. 8, pp. 843–851, 2002.

[70] S. Senanayake, D. Looi, M. Liew, K. Wong, and J. Hee, "Smart floor design for human gait analysis," *The Impact of Technology on Sport II*, p. 187, 2007.

[71] P. Srinivasan, D. Birchfield, G. Qian, and A. Kidané, "A pressure sensing floor for interactive media applications," in *ACM SIGCHI International Conference on Advances in Computer Entertainment Technology*. ACM, 2005, p. 281.

[72] C. Baker, K. Armijo, S. Belka, M. Benhabib, V. Bhargava, N. Burkhart, A. Minassians, G. Dervisoglu, L. Gutnik, M. Haick *et al.*, "Wireless sensor networks for home health care," in *21st International Conference on Advanced Information Networking and Applications Workshops (AINAW)*, vol. 2. Citeseer, 2007, pp. 832–837.

[73] K. Van Laerhoven, M. Borazio, D. Kilian, and B. Schiele, "Sustained logging and discrimination of sleep postures with low-level, wrist-worn sensors," in *12th IEEE International Symposium on Wearable Computers*. IEEE Computer Society, 2008, pp. 69–76.

[74] C. Hsia, Y. Hung, Y. Chiu, and C. Kang, "Bayesian classification for bed posture detection based on kurtosis and skewness estimation," in *10th International Conference on e-health Networking, Applications and Services*, 2008, pp. 165–168.

[75] K. Seo, C. Oh, T. Choi, and J. Lee, "Bed-type robotic system for the bedridden," in *2005 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, 2005, pp. 1170 1175.

[76] T. Harada, T. Mori, Y. Nishida, T. Yoshimi, and T. Sato, "Body parts positions and posture estimation system based on pressure distribution image," in *IEEE International Conference on Robotics and Automation*, vol. 2, 1999, pp. 968–975.

[77] Y. Nishida, M. Takeda, T. Mori, H. Mizoguchi, and T. Sato, "Monitoring patient respiration and posture using human symbiosis system," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'97)*, vol. 2, 1997.

[78] A. Arcelus, F. Knoefel, H. Sveistrup, and M. Bilodeau, "Determination of sit-to-stand transfer duration using bed and floor pressure sequences," *IEEE Transactions on Biomedical engineering*, vol. 56, no. 10, pp. 2485–2492, October 2009.

[79] B. Najafi, K. Aminian, A. Paraschiv-Ionescu, F. Loew, C. Bula, and P. Robert, "Ambulatory system for human motion analysis using a kinematic sensor: monitoring of daily physical activity in the elderly," *IEEE Transactions on Biomedical Engineering*, vol. 50, no. 6, pp. 711–723, 2003.

[80] M. Mathie, A. Coster, N. Lovell, and B. Celler, "Detection of daily physical activities using a triaxial accelerometer," *Medical and Biological Engineering and Computing*, vol. 41, no. 3, pp. 296–301, 2003.

[81] F. Allen, E. Ambikairajah, N. Lovell, and B. Celler, "Classification of a known sequence of motions and postures from accelerometry data using adapted Gaussian mixture models," *Physiological Measurement*, vol. 27, no. 10, pp. 935–952, 2006.

[82] N. Bidargaddi, A. Sarela, J. Boyle, V. Cheung, M. Karunanithi, L. Klingbei, C. Yelland, and L. Gray, "Wavelet based approach for posture transition esti-

mation using a waist worn accelerometer," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2007, pp. 1884–1887.

[83] G. Costantini, M. Carota, G. Maccioni, and D. Giansanti, "Classification of sit-to-stand locomotion task based on spectral analysis of waveforms generated by accelerometric transducer," *Electronics Letters*, vol. 42, p. 147, 2006.

[84] A. Salarian, H. Russmann, F. Vingerhoets, P. Burkhard, and K. Aminian, "Ambulatory monitoring of physical activities in patients with parkinson's disease," *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 12, pp. 2296–2299, 2007.

[85] M. Goffredo, M. Schmid, S. Conforto, M. Carli, A. Neri, and T. D'Alessio, "Markerless human motion analysis in Gauss–Laguerre transform domain: An application to sit-to-stand in young and elderly people," *IEEE Transactions on Information Technology in Biomedicine*, vol. 13, no. 2, pp. 209–216, March 2008.

[86] P. Millington, B. Myklebust, and G. Shambes, "Biomechanical analysis of the sit-to-stand motion in elderly persons," *Archives of Physical Medicine and Rehabilitation*, vol. 73, no. 7, p. 609, 1992.

[87] P. Cheng, M. Liaw, M. Wong, F. Tang, M. Lee, and P. Lin, "The sit-to-stand movement in stroke patients and its correlation with falling," *Archives of Physical Medicine and Rehabilitation*, vol. 79, no. 9, pp. 1043–1046, 1998.

[88] P. Riley, D. Krebs, and R. Popat, "Biomechanical analysis of failed sit-to-stand," *IEEE Transactions on Rehabilitation Engineering*, vol. 5, no. 4, pp. 353–359, 1997.

[89] M. Schenkman, R. Berger, P. Riley, R. Mann, and W. Hodge, "Whole-body

movements during rising to standing from sitting," *Physical Therapy*, vol. 70, no. 10, p. 638, 1990.

[90] A. Yang, S. Iyengar, S. Sastry, R. Bajcsy, P. Kuryloski, and R. Jafari, "Distributed segmentation and classification of human actions using a wearable motion sensor network," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. Citeseer, 2008, pp. 1–8.

[91] E. Guenterberg, H. Ghasemzadeh, R. Jafari, and R. Bajcsy, "A segmentation technique based on standard deviation in body sensor networks," in *IEEE Dallas Engineering in Medicine and Biology Workshop*, 2007.

[92] "Large sensor array," Tactex Controls Inc. Website, available at: http://www.tactex.com/sensor_array.

[93] "Kinotex," Tactex Controls Inc. Website, available at: http://www.tactex.com/kinotex.php.

[94] "Pressureguard renew," Span America Website, available at: http://www.spanamerica.com/renew.php.

[95] "Orange," Orange Website, software available at http://www.ailab.si/orange/.

[96] G. van Rossum et al., Python Language Website, available at http://www.python.org/.

[97] C.-C. Chang and C.-J. Lin, *LIBSVM: a library for support vector machines*, 2001, software available at http://www.csie.ntu.edu.tw/ cjlin/libsvm.

[98] R. Kohavi, P. Langley, and Y. Yun, "The utility of feature weighting in nearest-neighbor algorithms," in *9th European Conference on Machine Learning*. Prague: Springer-Verlag, 1997.

[99] A. Blum and P. Langley, "Selection of relevant features and examples in machine learning," *Artificial Intelligence*, vol. 97, no. 1-2, pp. 245–271, 1997.

[100] G. John, R. Kohavi, and K. Pfleger, "Irrelevant features and the subset selection problem," in *11th International Conference on Machine Learning*, vol. 129. Citeseer, 1994.

[101] P. R. Cohen, *Empirical Methods for Artificial Intelligence*. MIT Press, 1995.